JUNOS™ Internet Software Operational Mode Command Reference

Release 5.4

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
408-745-2000
www.juniper.net

Part Number: 530-007541-01, Revision 1

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986–1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by The Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, The Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., Copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks is registered in the U.S. Patent and Trademark Office and in other countries as a trademark of Juniper Networks, Inc. Broadband Cable Processor, G10, Internet Processor, JUNOS, JUNOScript, M5, M10, M20, M40, M40e, M160, M-series, T320, T640, and T-series are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

JUNOS Internet Software Chassis, Management, Policy, and Security Command Reference, Release 5.4
Copyright © 2002, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writer: Eva Moore
Editor: Stella Hackell
Covers and template design: Edmonds Design

Revision History
8 July 2002—First edition.

The information in this document is current as of the date listed in the revision history.

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.

Products made or sold by Juniper Networks (including the G10 CMTS, the M5 router, the M10 router, the M20 router, the M40 router, the M40e router, the M160 router, the T320 router, the T640 routing node, and the JUNOS software) or components thereof may be covered by one or more of the following patents which are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440.

YEAR 2000 NOTICE

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

SOFTWARE LICENSE

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

Abbreviated Table of Contents

About This Manual

Part 1
Overview

Chapter 1
Command-Line Interface Overview 3

Chapter 2
Strategies for Monitoring and Troubleshooting the Router 31

Part 2
System Management

Chapter 3
Monitor and Perform System Management Functions 41

Chapter 4
SNMP Monitoring and Troubleshooting 113

Chapter 5
M-Series Chassis Monitoring and Troubleshooting 123

Chapter 6
T640 Internet Routing Node Monitoring and Troubleshooting 161

Chapter 7
Accounting Options Monitoring and Troubleshooting 169

Part 3
Interfaces

Chapter 8
Router Interfaces Monitoring and Troubleshooting 179

Part 4
Policy Framework

Chapter 9
Firewall Filters Monitoring and Troubleshooting 477

Chapter 10
Forwarding Options Monitoring and Troubleshooting 481

Chapter 11
Routing Policy Monitoring and Troubleshooting 485

Part 5
Class of Service

Chapter 12
Class of Service Monitoring and Troubleshooting 491

Part 6
Protocols

Chapter 13
Protocol-Independent Routing Monitoring and Troubleshooting 515

Chapter 14
IS-IS Monitoring and Troubleshooting 569

Chapter 15
OSPF Monitoring and Troubleshooting 589

Chapter 16 BGP Monitoring and Troubleshooting 609

Chapter 17 RIP Monitoring and Troubleshooting 621

Chapter 18 IP Multicast Monitoring and Troubleshooting 625

Part 7 MPLS Applications

Chapter 19 MPLS Monitoring and Troubleshooting 675

Chapter 20 RSVP Monitoring and Troubleshooting 699

Chapter 21 LDP Monitoring and Troubleshooting 713

Chapter 22 Layer 2 VPN Monitoring and Troubleshooting 727

Chapter 23 Layer 3 VPN Monitoring and Troubleshooting 755

Part 8 Security

Chapter 24 IPSec Monitoring 793

Part 9
Appendix

Appendix A
Glossary 803

Part 10
Index

Index
Index 825

Index
Index of Statements and Commands 833

Table of Contents

About This Manual

Objectives xxxi
Audience xxxii
Document Organization xxxii
Chapter Organization xxxvi
Using the Indexes xxxvii
Documentation Conventions xxxvii
General Conventions xxxvii
Conventions for Software Commands and Statements xxxviii
List of Technical Publications xxxix
Documentation Feedback xl
How to Request Support xl

Part 1
Overview

Chapter 1
Command-Line Interface Overview 3

CLI Modes 3
CLI Command Hierarchy 4
CLI Operational Commands 5
Use the CLI 6
Get Help About Commands 6
Examples: Get Help with Commands 7
Have the CLI Complete Commands 8
Examples: Use CLI Command Completion 8
CLI Messages 9
Move Around and Edit the Command Line 10
How Output Appears on the Screen 10
Display Output One Screen at a Time 11
Filter Command Output 12
Place Command Output in a File 12
Search for a String in the Output 13
Count the Number of Lines in the Output 15
Display All Output at Once 15
Retain the Output after the Last Screen 15
Filter Command Output Multiple Times 16

Control the CLI Environment 16
Set the Terminal Type 16
Set the Screen Length 16
Set the Screen Width 16
Set the CLI Prompt 17
Idle Timeout 17
Set CLI to Prompt after a Software Upgrade 17
Set Command Completion 17
Display CLI Settings 18
Example: Control the CLI Environment 18
Set the Current Date and Time 18
Set the Current Date and Time from NTP Servers 18
Display CLI Command History 19
Monitor Who Uses the CLI 19
Summary of CLI Operational Commands 20
clear 20
configure 20
file 20
help 20
monitor 21
mtrace 21
ping 21
quit 21
| (pipe) 22
request 22
restart 23
set 23
show 23
ssh 24
start 24
telnet 24
test 24
traceroute 24
Summary of CLI Environment Commands 25
set cli complete-on-space 25
set cli idle-timeout 25
set cli prompt 26
set cli restart-on-upgrade 26
set cli screen-length 26
set cli screen-width 27
set cli terminal 27
set date 27
set date ntp 28
show cli 28
show cli history 29

Chapter 2
Strategies for Monitoring and Troubleshooting the Router 31

Basic Approaches to Troubleshooting 31
Troubleshooting Process 31
Identify the Symptoms 32
Isolate the Cause 32
Take Corrective Action 33
Evaluate the Solution 33
Example: Strategy for Isolating a Broken Network Connection 34
Tools for Troubleshooting 36
Commonly Used Operational Mode Commands 36
Craft Interface 37

Part 2
System Management

Chapter 3
Monitor and Perform System Management Functions 41

How to Specify Filenames and URLs 43
clear arp 44
clear log 44
clear system reboot 44
configure exclusive 44
file copy 45
file delete 45
file list 46
file rename 46
file show 46
monitor 47
ping 48
ping atm 50
ping vpn-interface 51
quit 51
request message 52
request support information 52
request system halt 53
request system logout user 54
request system reboot 55
request system snapshot 57
request system software add 59
request system software delete 61
request system software rollback 62
request system software validate 63
restart 64
show arp 65
show cli authorization 65
show cli history 66
show configuration 67

show host 68
show log 69
show ntp associations 70
show ntp status 71
show pfe fpc 72
show pfe next-hop 73
show pfe route 74
show pfe scb 75
show pfe sfm 77
show pfe statistics dma 79
show pfe statistics error 80
show pfe statistics ip 81
show pfe statistics notification 82
show pfe statistics pio 82
show pfe statistics traffic 83
show pfe terse 84
show system audit 85
show system boot-messages 86
show system buffers 88
show system connections 89
show system processes 91
show system reboot 97
show system software 97
show system statistics 98
show system storage 100
show system uptime 101
show system users 102
show system virtual-memory 103
show task 107
show version 109
ssh 110
start 111
telnet 112
test configuration 113
traceroute 114

Chapter 4
SNMP Monitoring and Troubleshooting 117

clear snmp statistics 118
show snmp rmon 119
show snmp statistics 123

Chapter 5 M-Series Chassis Monitoring and Troubleshooting 127

request chassis fpc slot 128
request chassis pcg slot 128
request chassis pic 129
request chassis routing-engine master 130
request chassis sfm slot 131
request chassis sfm master switch 132
request chassis ssb master switch 133
set chassis display message 134
show chassis alarms 135
show chassis craft-interface 136
show chassis environment 138
show chassis environment fpc 141
show chassis environment fpm 142
show chassis environment mcs 143
show chassis environment pcg 144
show chassis environment pem 145
show chassis environment routing-engine 146
show chassis environment sfm 147
show chassis feb 149
show chassis firmware 150
show chassis fpc 151
show chassis hardware 153
show chassis mac-addresses 157
show chassis routing-engine 158
show chassis scb 162
show chassis sfm 163
show chassis ssb 165

Chapter 6 T-Series Platform Monitoring and Troubleshooting 167

request chassis scg 167
request chassis sib 168
request chassis spmb restart slot 168
show chassis environment 169
show chassis environment cb 171
show chassis environment pem 172
show chassis environment scg 173
show chassis environment sib 174
show chassis spmb 174
show chassis spmb sibs 175

Chapter 7 Accounting Options Monitoring and Troubleshooting 177

show accounting profile 178
show accounting records 182

Part 3
Interfaces

Chapter 8
Static Router Interfaces Monitoring and Troubleshooting 187

Wildcard Characters in Interface Names 187
Examples: Use Wildcard Characters in Interface Names 188
Common Output Fields 189
Enabled Field 190
Device Flags Field 190
Link Flags Field 191
Interface Flags Field 191
Logical Interface Flags Field 192
Family Flags Field 192
Address Flags Field 192
clear interfaces statistics 193
show interfaces 194
show interfaces descriptions 195
show interfaces (filters | policers) 196
show interfaces media 199
show interfaces routing 200
show interfaces statistics 205
show interfaces terse 206

Chapter 9
ATM Interfaces Monitoring and Troubleshooting 209

show interfaces (for ATM Interfaces) 209
show interfaces (standard) (for ATM Interfaces) 217
show interfaces brief (for ATM Interfaces) 218
show interfaces detail (for ATM Interfaces) 218
show interfaces extensive (for ATM Interfaces) 219
show interfaces media (for ATM Interfaces) 222
show interfaces statistics (for ATM Interfaces) 222

Chapter 10
Channelized DS-3 to DS-0 Interfaces Monitoring and Troubleshooting 225

show interfaces (for Channelized DS-3 to DS-0 Interfaces) 225
show interfaces (standard) (for Channelized DS-3 to DS-0 Interfaces) 233
show interfaces brief (for Channelized DS-3 to DS-0 Interfaces) 233
show interfaces detail (for Channelized DS-3 to DS-0 Interfaces) 234
show interfaces extensive (for Channelized DS-3 to DS-0 Interfaces) 235
show interfaces media (for Channelized DS-3 to DS-0 Interfaces) 237
show interfaces statistics (for Channelized DS-3 to DS-0 Interfaces) 237

Chapter 11
Channelized DS-3 to DS-1 Interfaces Monitoring and Troubleshooting 239

show interfaces (for Channelized DS-3 to DS-1 Interfaces) 239
show interfaces (standard) for Channelized DS-3 to DS-1 Interfaces) 246
show interfaces brief for Channelized DS-3 to DS-1 Interfaces) 246
show interfaces detail for Channelized DS-3 to DS-1 Interfaces) 247
show interfaces extensive for Channelized DS-3 to DS-1 Interfaces) 248
show interfaces media for Channelized DS-3 to DS-1 Interfaces) 250
show interfaces statistics for Channelized DS-3 to DS-1 Interfaces) 251

Chapter 12
Channelized E-1 Interfaces Monitoring and Troubleshooting 253

show interfaces (for Channelized E1 Interfaces) 253
show interfaces (standard) (for Channelized E1 interfaces) 261
show interfaces brief (for Channelized E1 interfaces) 262
show interfaces detail (for Channelized E1 interfaces) 262
show interfaces extensive (for Channelized E1 interfaces) 264
show interfaces media (for Channelized E1 interfaces) 266
show interfaces statistics (for Channelized E1 interfaces) 266

Chapter 13
Channelized OC-12 Interfaces Monitoring and Troubleshooting 269

show interfaces (for Channelized OC-12 Interfaces) 269
show interfaces (standard) (for Channelized OC-12 Interfaces) 277
show interfaces brief (for Channelized OC-12 Interfaces) 278
show interfaces detail (for Channelized OC-12 Interfaces) 278
show interfaces extensive (for Channelized OC-12 Interfaces) 279
show interfaces media (for Channelized OC-12 Interfaces) 281
show interfaces statistics (for Channelized OC-12 Interfaces) 282

Chapter 14
Channelized STM-1 Interfaces Monitoring and Troubleshooting 283

show interfaces (for Channelized STM-1 Interfaces) 283
show interfaces (standard) (for Channelized STM-1 Interfaces) 292
show interfaces brief (for Channelized STM-1 Interfaces) 293
show interfaces detail (for Channelized STM-1 Interfaces) 293
show interfaces extensive (for Channelized STM-1 Interfaces) 294
show interfaces media (for Channelized STM-1 Interfaces) 296
show interfaces statistics (for Channelized STM-1 Interfaces) 297

Chapter 15
E-1 Interfaces Monitoring and Troubleshooting 299

show interfaces (for E1 Interfaces) 299
show interfaces (standard) (for E1 interfaces) 307
show interfaces brief (for E1 interfaces) 308
show interfaces detail (for E1 interfaces) 308
show interfaces extensive (for E1 interfaces) 309
show interfaces media (for E1 interfaces) 311
show interfaces statistics (for E1 interfaces) 312

Chapter 16
E-3 Interfaces Monitoring and Troubleshooting 313

show interfaces (for E3 Interfaces) 313
show interfaces (standard) (for E3 interfaces) 320
show interfaces brief (for E3 interfaces) 320
show interfaces detail (for E3 interfaces) 320
show interfaces extensive (for E3 interfaces) 321
show interfaces media (for E3 interfaces) 322
show interfaces statistics (for E3 interfaces) 323

Chapter 17
ES Interfaces Monitoring and Troubleshooting 325

show interfaces (for ES Interfaces) 325
show interfaces (standard) (for ES interfaces) 330
show interfaces brief (for ES interfaces) 330
show interfaces detail (for ES interfaces) 331
show interfaces extensive (for ES interfaces) 332
show interfaces media (for ES interfaces) 332
show interfaces statistics (for ES interfaces) 333

Chapter 18
Fast Ethernet Interfaces Monitoring and Troubleshooting 335

show interfaces (for Fast Ethernet Interfaces) 335
show interfaces (standard) (for Fast Ethernet interfaces) 342
show interfaces brief (for Fast Ethernet interfaces) 342
show interfaces detail (for Fast Ethernet interfaces) 343
show interfaces extensive (for Fast Ethernet interfaces) 344
show interfaces media (for Fast Ethernet interfaces) 345
show interfaces statistics (for Fast Ethernet interfaces) 346

Chapter 19
Gigabit Ethernet Interfaces Monitoring and Troubleshooting 347

show interfaces (for Gigabit Ethernet Interfaces) 347
show interfaces (standard) (for Gigabit Ethernet interfaces) 354
show interfaces (standard) (for Aggregated Gigabit Ethernet interfaces) 354
show interfaces brief (for Gigabit Ethernet interfaces) 355
show interfaces detail (for Gigabit Ethernet interfaces) 355
show interfaces detail (for Aggregated Gigabit Ethernet interfaces) 356
show interfaces extensive (for Gigabit Ethernet interfaces) 357
show interfaces extensive (for Aggregated Gigabit Ethernet interfaces) 358
show interfaces media (for Gigabit Ethernet interfaces) 359
show interfaces statistics (for Gigabit Ethernet interfaces) 359

Chapter 20
GRE Interfaces Monitoring and Troubleshooting 361

show interfaces (for GRE Interfaces) 361
show interfaces (standard) (for GRE Interfaces) 366
show interfaces brief (for GRE Interfaces) 366
show interfaces detail (for GRE Interfaces) 367
show interfaces extensive (for GRE Interfaces) 368
show interfaces media (for GRE Interfaces) 369
show interfaces statistics (for GRE Interfaces) 369

Chapter 21
IP-IP Interfaces Monitoring and Troubleshooting 371

show interfaces (for IP-IP Interfaces) 371
show interfaces (standard) (for IP-IP interfaces) 376
show interfaces brief (for IP-IP interfaces) 377
show interfaces detail (for IP-IP interfaces) 377
show interfaces extensive (for IP-IP interfaces) 378
show interfaces media (for IP-IP interfaces) 378
show interfaces statistics (for IP-IP interfaces) 379

Chapter 22
Multilink Services Interfaces Monitoring and Troubleshooting 381

show interfaces (for Multilink Services Interfaces) 381
show interfaces (standard) (for Multilink Services Interfaces) 387
show interfaces brief (for Multilink Services Interfaces) 388
show interfaces detail (for Multilink Services Interfaces) 388
show interfaces extensive (for Multilink Services Interfaces) 389
show interfaces media (for Multilink Services Interfaces) 391
show interfaces statistics (for Multilink Services Interfaces) 391

Chapter 23
SONET/SDH Interfaces Monitoring and Troubleshooting 393

show interfaces (for SONET/SDH Interfaces) 393
show interfaces (standard) (for SDH/SONET interfaces) 402
show interfaces brief (for SDH/SONET interfaces) 403
show interfaces detail (for SDH/SONET interfaces) 403
show interfaces detail (for Aggregated SDH/SONET interfaces) 404
show interfaces extensive (for SDH/SONET interfaces) 405
show interfaces extensive (for Aggregated SDH/SONET interfaces) 407
show interfaces media (for SDH/SONET interfaces) 408
show interfaces statistics (for SDH/SONET interfaces) 408

Chapter 24
T1 Interfaces Monitoring and Troubleshooting 409

show interfaces (for T1 Interfaces) 409
show interfaces (standard) (for T1 Interfaces) 417
show interfaces brief (for T1 Interfaces) 417
show interfaces detail (for T1 Interfaces) 417
show interfaces extensive (for T1 Interfaces) 418
show interfaces media (for T1 Interfaces) 419
show interfaces statistics (for T1 Interfaces) 419

Chapter 25
T3 Interfaces Monitoring and Troubleshooting 421

show interfaces (for T3 Interfaces) 421
show interfaces (standard) (for T3 interfaces) 429
show interfaces brief (for T3 interfaces) 429
show interfaces detail (for T3 interfaces) 430
show interfaces extensive (for T3 interfaces) 431
show interfaces media (for T3 interfaces) 432
show interfaces statistics (for T3 interfaces) 433

Chapter 26
Virtual Tunnel Interfaces Monitoring and Troubleshooting 435

show interfaces (for VT Interfaces) 435
show interfaces (standard) (for VT interfaces) 441
show interfaces brief (for VT interfaces) 441
show interfaces detail (for VT interfaces) 442
show interfaces extensive (for VT interfaces) 442
show interfaces media (for VT interfaces) 443
show interfaces statistics (for VT interfaces) 443

Chapter 27
Real-Time Router Interface Monitoring and Troubleshooting 445

Monitor Traffic Match Conditions 445
monitor interface 449
monitor traffic 452
show aps 455

Chapter 28
Monitoring and Troubleshooting using BERT Tests 459

test interface bert-start 459
test interface bert-stop 460
test interface fdl-line-loop 461
test interface fdl-payload-loop 462
test interface feac-loop-initiate 463
test interface feac-loop-terminate 463

Chapter 29
VRRP Router Interfaces Monitoring and Troubleshooting 465

clear vrrp 465
show vrrp 466

Part 4
Policy Framework

Chapter 30
Firewall Filters Monitoring and Troubleshooting 475

clear firewall 475
show firewall filter 476
show firewall log 477

Chapter 31
Forwarding Options Monitoring and Troubleshooting 479

clear helper statistics 479
show helper statistics 480
restart helper 481

Chapter 32
Routing Policy Monitoring and Troubleshooting 483

show policy 483
show policy damping 484
test policy 484

Part 5
Class of Service

Chapter 33
Class of Service Monitoring and Troubleshooting 487

show class-of-service 488
show class-of-service classifier 489
show class-of-service code-point-aliases 490
show class-of-service drop-profile 491
show class-of-service forwarding-class 492
show class-of-service forwarding-table 492
show class-of-service forwarding-table classifier 493
show class-of-service forwarding-table classifier mapping 494
show class-of-service forwarding-table drop-profile 495
show class-of-service forwarding-table rewrite-rule 496
show class-of-service forwarding-table rewrite-rule mapping 497
show class-of-service forwarding-table scheduler-map 498
show class-of-service interface 500
show class-of-service rewrite-rule 501

show class-of-service scheduler-map 502
show interfaces extensive (CoS related attributes only) 504
show interfaces queue 506

Part 6
Protocols

Chapter 34
Protocol-Independent Routing Monitoring and Troubleshooting

Protocol-Independent Routing Monitoring and Troubleshooting 511

clear route forwarding-table 513
show as-path 514
show route 516
show route advertising-protocol 527
show route aspath-regex 529
show route best 531
show route community 532
show route damping 534
show route exact 537
show route export 539
show route export vrf-target 541
show route forwarding-table 543
show route inactive 549
show route instance 550
show route label-switched-path 552
show route martians 554
show route next-hop 555
show route output 556
show route protocol 557
show route range 559
show route receive-protocol 561
show route resolution 563
show route source-gateway 565
show route summary 566
show route table 567
show route terse 571

Chapter 35
IS-IS Monitoring and Troubleshooting 573

System ID Format in Command Output 573
clear isis adjacency 574
clear isis database 575
clear isis statistics 576
show isis adjacency 577
show isis database 580
show isis hostname 583
show isis interface 584
show isis routes 587
show isis spf 588
show isis statistics 591

Chapter 36
OSPF Monitoring and Troubleshooting 593

clear ospf database 594
clear ospf neighbor 594
clear ospf statistics 595
show ospf database 596
show ospf interface 602
show ospf io-statistics 605
show ospf log 606
show ospf neighbor 607
show ospf route 610
show ospf statistics 612

Chapter 37
BGP Monitoring and Troubleshooting 613

clear bgp damping 613
clear bgp neighbor 614
show bgp group 615
show bgp neighbor 617
show bgp summary 622

Chapter 38
RIP Monitoring and Troubleshooting 625

clear rip statistics 625
show rip neighbor 626
show rip statistics 627

Chapter 39
RIPng Monitoring and Troubleshooting 629

clear ripng statistics 629
show ripng neighbor 630
show route protocol ripng table inet6 631
show ripng statistics 632

Chapter 40
IPv6 Monitoring and Troubleshooting 635

clear ipv6 neighbors 635
clear ipv6 router-advertisement 636
show ipv6 neighbors 636
show ipv6 router-advertisement 637

Chapter 41
IP Multicast Monitoring and Troubleshooting 639

clear igmp membership 641
clear igmp statistics 643
clear msdp cache 644
clear msdp statistics 644
clear multicast statistics 644
clear pim join 645
clear pim statistics 646
mtrace 647
mtrace from-source 648
mtrace monitor 650
mtrace to-gateway 651
show dvmrp interfaces 653
show dvmrp neighbors 654
show dvmrp prefix 655
show dvmrp prunes 657
show igmp group 658
show igmp interface 660
show igmp statistics 661
show msdp 662
show msdp statistics 663
show multicast mrinfo 664
show multicast nexthops 665
show multicast route 666
show multicast rpf 669
show multicast scope 671
show multicast sessions 672
show multicast statistics 673
show multicast tunnels 674
show multicast usage 675
show pim bootstrap 676
show pim interfaces 677
show pim join 679

show pim neighbors 682
show pim rps 683
show pim source 686
show pim statistics 687
show sap listen 688
test msdp 688

Part 7 MPLS Applications

Chapter 42 MPLS Monitoring and Troubleshooting 691

clear mpls lsp 692
show connections 693
show link-management 695
show mpls admin-groups 696
show mpls cspf 697
show mpls interface 698
show mpls lsp 699
show mpls path 709
show ted database 710
show ted link 714
show ted protocol 716

Chapter 43 RSVP Monitoring and Troubleshooting 717

clear rsvp session 717
clear rsvp statistics 718
show rsvp interface 719
show rsvp neighbor 724
show rsvp session 727
show rsvp statistics 734
show rsvp version 736

Chapter 44
LDP Monitoring and Troubleshooting....737

clear ldp neighbor 737
clear ldp statistics 737
show ldp database 738
show ldp interface 741
show ldp neighbor 742
show ldp path 743
show ldp route 745
show ldp session 747
show ldp statistics 749

Chapter 45
Layer 2 VPN Monitoring and Troubleshooting....751

show bgp neighbor 752
show bgp summary 757
show interfaces terse 759
show l2circuit connections 760
show l2vpn connections 765
show route 768
show route advertising-protocol 769
show route forwarding-table 771
show route receive-protocol 774
show route table 777

Chapter 46
Layer 3 VPN Monitoring and Troubleshooting....779

ping vpn-interface 780
show bgp neighbor 781
show bgp summary 787
show ospf database 789
show ospf interface 790
show ospf log 790
show ospf neighbor 791
show ospf route 791
show ospf statistics 792
show rip neighbor 793
show rip statistics 793
show route 794
show route advertising-protocol 795
show route forwarding-table 797
show route instance 800
show route protocol 802
show route receive-protocol 805
show route table 809
ssh 812
telnet 813
traceroute 814

Part 8
Security

Chapter 47
IPSec Monitoring 817

clear ike security-associations 817
clear ipsec security-associations 818
show ike security-associations 819
show ipsec security-associations 822

Part 9
Appendix

Appendix A
Glossary 827

Part 10
Index

Index
Index 849

Index
Index of Statements and Commands 861

List of Figures

List of Figures

Figure 1: CLI Command Hierarchy Example....4
Figure 2: Troubleshooting Process....32
Figure 3: Example: Troubleshooting Loss of Connectivity....35

List of Tables

List of Tables

Table 1: Juniper Networks Technical Documentation xxxix
Table 2: CLI Keyboard Sequences 10
Table 3: ---More--- Prompt Keyboard Sequences 11
Table 4: Common Regular Expression Operators 13
Table 5: Commonly Used Operational Mode Commands 36
Table 6: Commands for Monitoring and Performing System Management 41
Table 7: Show System Processes Output Field Summary 92
Table 8: Commands for Monitoring and Managing SNMP 117
Table 9: Show SNMP Rmon Output Field Summary 119
Table 10: Commands for Monitoring the M-series Chassis 127
Table 11: Routing Engine Mastership Election 130
Table 12: Commands for Monitoring the T-series platforms 167
Table 13: Commands for Monitoring and Managing Accounting Options 177
Table 14: Commands for Monitoring Interfaces 187
Table 15: Wildcard Characters for Specifying Interface Names 188
Table 16: Show Interfaces Routing Output Field Summary 201
Table 17: ATM Show Interfaces Output Field Summary in Alphabetical Order 210
Table 18: ATM Show Interfaces Output Field Summary in Order of Appearance 212
Table 19: Channelized DS-3 to DS-0 Show Interfaces Output Field Summary in Alphabetical Order 226
Table 20: Channelized DS-3 to DS-0 Show Interfaces Output Field Summary in Order of Appearance 228
Table 21: Channelized DS-3 to DS-1 Show Interfaces Output Field Summary 240
Table 22: Channelized DS-3 to DS-1 Show Interfaces Output Field Summary in Order of Appearance 242
Table 23: Channelized E1 Show Interfaces Output Field Summary 254
Table 24: Channelized E1 Show Interfaces Output Field Summary in Order of Appearance 256
Table 25: Channelized OC-12 Show Interfaces Output Field Summary 270
Table 26: Channelized OC-12 Show Interfaces Output Field Summary in Order of Appearance 272
Table 27: Channelized STM-1 Interfaces Output Field Summary 284
Table 28: Channelized STM-1 Show Interfaces Output Field Summary in Order of Appearance 286
Table 29: E1 Show Interfaces Output Field Summary 300
Table 30: E1 Show Interfaces Output Field Summary in Order of Appearance 302
Table 31: E3 Show Interfaces Output Field Summary 314
Table 32: E3 Show Interfaces Output Field Summary in Order of Appearance 315
Table 33: ES Show Interfaces Output Field Summary 326

Table 34: ES Show Interfaces Output Field Summary in Order of Appearance....328
Table 35: Fast Ethernet Show Interfaces Output Field Summary 336
Table 36: Fast Ethernet Show Interfaces Output Field Summary in Order of Appearance 338
Table 37: Gigabit Ethernet Show Interfaces Output Field Summary 348
Table 38: Gigabit Ethernet Show Interfaces Output Field Summary in Order of Appearance 350
Table 39: GRE Show Interfaces Output Field Summary 362
Table 40: GRE Show Interfaces Output Field Summary in Order of Appearance....364
Table 41: IP-IP Show Interfaces Output Field Summary 372
Table 42: IP-IP Show Interfaces Output Field Summary in Order of Appearance....374
Table 43: Multilink Services Show Interfaces Output Field Summary 382
Table 44: SONET/SDH Show Interfaces Output Field Summary 394
Table 45: SONET/SDH Show Interfaces Output Field Summary in Order of Appearance 396
Table 46: T1 Show Interfaces Output Field Summary 410
Table 47: T1 Show Interfaces Output Field Summary in Order of Appearance....412
Table 48: T3 Show Interfaces Output Field Summary 422
Table 49: T3 Show Interfaces Output Field Summary in Order of Appearance....424
Table 50: VT Show Interfaces Output Field Summary 436
Table 51: VT Show Interfaces Output Field Summary in Order of Appearance....438
Table 52: Commands for Real-Time Interface Monitoring and Troubleshooting....445
Table 53: Match Conditions for the monitor traffic Command 445
Table 54: Logical Operators for the monitor traffic Command 447
Table 55: Arithmetic and Relational Operators for the monitor traffic Command 448
Table 56: Monitor Interface Output Control Keys 449
Table 57: Show APS Output Field Summary....455
Table 58: Commands for Monitoring Interfaces Using BERT Tests 459
Table 59: Table 59 Commands for Monitoring VRRP Interfaces 465
Table 60: Show VRRP Output Field Summary 467
Table 61: Commands for Monitoring and Managing Firewall Filters 475
Table 62: Commands for Forwarding Options Monitoring and Troubleshooting....479
Table 63: Commands for Monitoring and Managing Routing Policy 483
Table 64: Commands for Monitoring and Troubleshooting Class of Service....487
Table 65: Commands for Monitoring Protocol-Independent Routing Properties 512
Table 66: Show Route Output Field Summary 517
Table 67: Commands for Monitoring IS-IS 573
Table 68: Show IS-IS Adjacency Output Field Summary 577
Table 69: Show ISIS Database Output Field Summary....580
Table 70: Show IS-IS Interface Output Field Summary 584
Table 71: Show ISIS SPF Interface Output Field Summary 588
Table 72: Commands for Monitoring OSPF 593
Table 73: Show OSPF Database Output Field Summary 597
Table 74: Show OSPF Interface Output Field Summary....602
Table 75: Show OSPF Neighbor Output Field Summary 607
Table 76: Show OSPF Route Output Fields Summary 610

Table 77: Commands for Monitoring BGP 613
Table 78: Commands for Monitoring RIP 625
Table 79: Commands for Monitoring RIPng 629
Table 80: Commands for Monitoring IPv6 635
Table 81: Commands for Monitoring IP Multicast 639
Table 82: Show DVMRP Prefix Output Field Summary 655
Table 83: Show Multicast Route Output Field Summary 667
Table 84: Show Multicast RPF Output Field Summary 669
Table 85: Show PIM Join Output Field Summary 679
Table 86: Show PIM RPS Output Fields Summary 683
Table 87: Commands for Monitoring MPLS 691
Table 88: Show MPLS LSP Output Field Summary 700
Table 89: Show TED Database Output Fields 710
Table 90: Show TED Link Output Field Summary 714
Table 91: Commands for Monitoring RSVP 717
Table 92: Show RSVP Interface Output Fields Summary 719
Table 93: Show RSVP Session Output Field Summary 728
Table 94: Commands for Monitoring LDP 737
Table 95: Show LDP Database Output Fields Summary 738
Table 96: Commands for Monitoring Layer 2 VPNs.....751
Table 97: Show l2circuit connections Output Field Summary 761
Table 98: Commands for Monitoring Layer 3 VPNs 779
Table 99: Commands for Monitoring Security 817

About This Manual

This chapter provides a high-level overview of the *JUNOS Internet Software Operational Mode Command Reference:*

- Objectives on page xxxi
- Audience on page xxxii
- Document Organization on page xxxii
- Chapter Organization on page xxxvi
- Using the Indexes on page xxxvii
- Documentation Conventions on page xxxvii
- List of Technical Publications on page xxxix
- Documentation Feedback on page xl
- How to Request Support on page xl

Objectives

This manual provides an overview of the JUNOS Internet software commands that you use to monitor and troubleshoot the router.

This manual documents Release 5.4 of the JUNOS Internet software. To obtain additional information about the JUNOS software—either corrections to information in this manual or information that might have been omitted from this manual—refer to the software release notes.

To obtain the most current version of this manual and the most current version of the software release notes, refer to the product documentation page on the Juniper Networks Web site, which is located at http://www.juniper.net/.

To order printed copies of this manual or to order a documentation CD-ROM, which contains this manual, please contact your sales representative.

Audience

This manual is designed for network administrators who are monitoring a Juniper Networks router. It assumes that you have a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. This manual assumes that you are familiar with one or more of the following Internet routing protocols: Border Gateway Protocol (BGP), Routing Information Protocol (RIP), Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), Internet Control Message Protocol (ICMP) router discovery, Internet Group Management Protocol (IGMP), Distance Vector Multicast Routing Protocol (DVMRP), Protocol-Independent Multicast (PIM), Multiprotocol Label Switching (MPLS), Resource Reservation Protocol (RSVP), and Simple Network Management Protocol (SNMP).

Document Organization

This manual is divided into several parts. Each part describes a major functional area of the JUNOS software, and the individual chapters within a part describe the software commands of that functional area.

This manual contains the following parts and chapters:

- Preface, “About This Manual” (this chapter), provides a brief description of the contents and organization of this manual and describes how to contact customer support.
- Part 1, “Overview,” provides an overview of the *JUNOS Internet Software Operational Mode Command Reference*, describes the user command-line interface, and provides basic strategies for monitoring and troubleshooting the router:
  - Chapter 1, “Command-Line Interface Overview,” describes the basics of the interface that you use to monitor the JUNOS software. The CLI is the interface that you use whenever you access the router. For information about configuring the router, or for additional information about the CLI, see the *JUNOS Internet Software Configuration Guide*: *Getting Started.*
  - Chapter 2, “Strategies for Monitoring and Troubleshooting the Router,” describes basic strategies for monitoring and troubleshooting the network and lists the commands most commonly used for monitoring and troubleshooting.
- Part 2, “System Management,” describes the CLI commands used to manage the router and chassis:
  - Chapter 3, “Monitor and Perform System Management Functions,” describes the commands for monitoring user access and performing system management.
  - Chapter 4, “SNMP Monitoring and Troubleshooting,” describes the commands for monitoring user access and performing SNMP management.
  - Chapter 5, “M-Series Chassis Monitoring and Troubleshooting,” describes the commands for monitoring the router chassis and troubleshooting chassis problems.

- Chapter 6, “T-Series Platform Monitoring and Troubleshooting,” describes the commands for monitoring the T640 Internet routing node chassis and troubleshooting chassis problems.
- Chapter 7, “Accounting Options Monitoring and Troubleshooting,” describes the commands for monitoring accounting options and troubleshooting accounting options problems.

- Part 3, “Interfaces,” describes the CLI commands used to manage the router interfaces:
  - Chapter 8, “Static Router Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring static interfaces and troubleshooting static interface problems.
  - Chapter 9, “ATM Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring ATM interfaces and troubleshooting ATM interface problems.
  - Chapter 10, “Channelized DS-3 to DS-0 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Channelized DS-3 to DS-0 router interfaces and troubleshooting Channelized DS-3 to DS-0 router interface problems.
  - Chapter 11, “Channelized DS-3 to DS-1 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Channelized DS-3 to DS-1 router interfaces and troubleshooting Channelized DS-3 to DS-1 router interface problems.
  - Chapter 12, “Channelized E-1 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Channelized E-1 router interfaces and troubleshooting Channelized E-1 router interface problems.
  - Chapter 13, “Channelized OC-12 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Channelized OC-12 router interfaces and troubleshooting Channelized OC-12 router interface problems.
  - Chapter 14, “Channelized STM-1 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Channelized STM-1 router interfaces and troubleshooting Channelized STM-1 router interface problems.
  - Chapter 15, “E-1 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring E-1 router interfaces and troubleshooting E-1 router interface problems.
  - Chapter 16, “E-3 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring E-3 router interfaces and troubleshooting E-3 router interface problems.
  - Chapter 17, “ES Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring ES router interfaces and troubleshooting ES router interface problems.

- Chapter 18, “Fast Ethernet Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring ES router interfaces and troubleshooting ES router interface problems.
- Chapter 19, “Gigabit Ethernet Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Gigabit Ethernet router interfaces and troubleshooting Gigabit Ethernet router interface problems.
- Chapter 20, “GRE Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring GRE router interfaces and troubleshooting GRE router interface problems.
- Chapter 21, “IP-IP Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring IP-IP router interfaces and troubleshooting IP-IP router interface problems.
- Chapter 22, “Multilink Services Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Multilink Services router interfaces and troubleshooting Multilink Services router interface problems.
- Chapter 23, “SONET/SDH Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring SONET/SDH router interfaces and troubleshooting SONET/SDH router interface problems.
- Chapter 24, “T1 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring T1 router interfaces and troubleshooting T1router interface problems.
- Chapter 25, “T3 Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring T3 router interfaces and troubleshooting T3 router interface problems.
- Chapter 26, “Virtual Tunnel Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring Virtual Tunnel router interfaces and troubleshooting Virtual Tunnel router interface problems.
- Chapter 27, “Real-Time Router Interface Monitoring and Troubleshooting,” describes the commands for monitoring real-time router interfaces and troubleshooting real-time router interface problems.
- Chapter 28, “Monitoring and Troubleshooting using BERT Tests,” describes the commands for monitoring BERT tests and troubleshooting problems.
- Chapter 29, “VRRP Router Interfaces Monitoring and Troubleshooting,” describes the commands for monitoring VRRP router interfaces and troubleshooting VRRP router interface problems.

- Part 4, “Policy Framework,” describes the CLI commands used to manage firewall filters and routing policy:
  - Chapter 30, “Firewall Filters Monitoring and Troubleshooting,” describes the commands for monitoring firewall filters and troubleshooting filter problems.
  - Chapter 31, “Forwarding Options Monitoring and Troubleshooting,” describes the commands for monitoring forwarding options and troubleshooting forwarding option problems.
  - Chapter 32, “Routing Policy Monitoring and Troubleshooting,” describes the commands for monitoring routing policy and troubleshooting routing policy problems.
- Part 5, “Class of Service,” describes the CLI commands used to manage Class of Service:
  - Chapter 33, “Class of Service Monitoring and Troubleshooting,” describes the commands for monitoring Class of Service (CoS) and troubleshooting CoS problems.
- Part 6, “Protocols,” describes the CLI commands used to monitor the JUNOS software routing policy, unicast routing protocols, and multicast protocols:
  - Chapter 34, “Protocol-Independent Routing Monitoring and Troubleshooting,” describes the commands for monitoring and troubleshooting protocol-independent routing properties.
  - Chapter 35, “IS-IS Monitoring and Troubleshooting,” describes the commands for monitoring IS-IS and troubleshooting IS-IS problems.
  - Chapter 36, “OSPF Monitoring and Troubleshooting,” describes the commands for monitoring OSPF and troubleshooting OSPF problems.
  - Chapter 37, “BGP Monitoring and Troubleshooting,” describes the commands for monitoring BGP and troubleshooting BGP problems.
  - Chapter 38, “RIP Monitoring and Troubleshooting,” describes the commands for monitoring RIP and troubleshooting RIP problems.
  - Chapter 39, “RIPng Monitoring and Troubleshooting,” describes the commands for monitoring RIPng and troubleshooting RIPng problems.
  - Chapter 40, “IPv6 Monitoring and Troubleshooting,” describes the commands for monitoring IPv6 and troubleshooting IPv6 problems.
  - Chapter 41, “IP Multicast Monitoring and Troubleshooting,” describes the CLI commands for monitoring multicast protocols and troubleshooting multicast problems.

- Part 7, “MPLS Applications,” describes the CLI commands used to manage traffic engineering:
  - Chapter 42, “MPLS Monitoring and Troubleshooting,” describes the commands for monitoring MPLS and troubleshooting MPLS problems.
  - Chapter 43, “RSVP Monitoring and Troubleshooting,” describes the commands for monitoring RSVP and troubleshooting RSVP problems.
  - Chapter 44, “LDP Monitoring and Troubleshooting,” describes the commands for monitoring LDP and troubleshooting LDP problems.
  - Chapter 45, “Layer 2 VPN Monitoring and Troubleshooting,” describes the commands for monitoring Layer 2 VPNs and troubleshooting Layer 2 VPN problems.
  - Chapter 46, “Layer 3 VPN Monitoring and Troubleshooting,” describes the commands for monitoring Layer 3 VPNs and troubleshooting Layer 3 VPN problems.
- Part 8, “Security”, describes the commands used to monitor security:
  - Chapter 47, “IPSec Monitoring”, describes the commands for monitoring and troubleshooting IPSec problems.

This manual also contains a glossary, a complete index, and an index of statements and commands.

Chapter Organization

Most chapters in this manual consist of a brief description of a specific software component followed by the commands that you use for monitoring and troubleshooting that component. A table at the beginning of the chapter lists commands available for monitoring and troubleshooting the software components. Then, the commands are explained alphabetically.

The commands in this manual typically contain the following information:

- Syntax—Describes the full syntax of the command. For an explanation of how to read the syntax statements, see “Documentation Conventions” on page xxxvii.
- Description—Describes the function of the command.
- Options—Describes the command’s options, if there are any. For options with numeric values, the allowed range and default value, if any, are listed. For multiple options, if one option is the default, that fact is stated.
- Required privilege level—Indicates the permissions that the user must have to view the command. For an explanation of the permissions, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

- See also—Indicates other commands or statements that might provide related or similar functionality.
- Output fields—Describes each possible output field for the command, including ranges and default values. These fields are first listed alphabetically, then listed in more detail in the order of appearance.
- Sample Output—Display typical sample output for the standard, brief, summary, terse, detail, and extensive options of a command (when applicable).

Using the Indexes

This manual contains two indexes: a complete index, which contains all index entries, and an index that contains only statements and commands.

In the complete index, bold page numbers point to pages in the statement summary chapters. The index entry for each configuration statement always contains at least two entries. The first, with a bold page number on the same line as the statement name, references the statement summary section. The second entry, "usage guidelines," references the section in a configuration guidelines chapter that describes how to use the statement.

Documentation Conventions

General Conventions

This manual uses the following text conventions:

- Statements, commands, filenames, directory names, IP addresses, and configuration hierarchy levels are shown in a sans serif font. In the following example, *stub* is a statement name and [edit protocols ospf area *area-id*] is a configuration hierarchy level:

  To configure a stub area, include the stub statement at the [edit protocols ospf area *area-id*] hierarchy level:

- In examples, text that you type literally is shown in bold. In the following example, you type the word *show*:

  [edit protocols ospf area *area-id*]
  cli# **show**
  stub <default-metric *metric*>

- Examples of command output are generally shown in a fixed-width font to preserve the column alignment. For example:

```
> show interfaces terse
Interface       Admin Link Proto Local             Remote
at-1/3/0        up    up
at-1/3/0.0      up    up   inet  1.0.0.1           --> 1.0.0.2
                           iso
fxp0            up    up
fxp0.0          up    up   inet  192.168.5.59/24
```

Conventions for Software Commands and Statements

When describing the JUNOS software, this manual uses the following type and presentation conventions:

- Statement or command names that you type literally are shown nonitalicized. In the following example, the statement name is *area*:

  You configure all these routers by including the following area statement at the [edit protocols ospf] hierarchy level:

- Options, which are variable terms for which you substitute appropriate values, are shown in italics. In the following example, *area-id* is an option. When you type the area statement, you substitute a value for *area-id*.

  area *area-id*;

- Optional portions of a configuration statement are enclosed in angle brackets. In the following example, the "default-metric *metric*" portion of the statement is optional:

  stub <default-metric *metric*>;

- For text strings separated by a pipe (|), you must specify either *string1* or *string2*, but you cannot specify both or neither of them. Parentheses are sometimes used to group the strings.

  string1 | *string2*
  (*string1* | *string2*)

  In the following example, you must specify either broadcast or multicast, but you cannot specify both:

  broadcast | multicast

- For some statements, you can specify a set of values. The set must be enclosed in square brackets. For example:

  community *name* members [*community-ids*]

- The configuration examples in this manual are generally formatted in the way that they appear when you issue a show command. This format includes braces ({ }) and semicolons. When you type configuration statements in the CLI, you do not type the braces and semicolons. However, when you type configuration statements in an ASCII file, you must include the braces and semicolons. For example:

```
[edit]
cli# set routing-options static route default nexthop address retain
[edit]
cli# show
routing-options {
    static {
        route default {
            nexthop address;
            retain;
        }
    }
}
```

- Comments in the configuration examples are shown either preceding the lines that the comments apply to, or more often, they appear on the same line. When comments appear on the same line, they are preceded by a pound sign (#) to indicate where the comment starts. In an actual configuration, comments can only precede a line; they cannot be on the same line as a configuration statement. For example:

```
protocols {
    mpls {
        interface (interface-name | all);    # Required to enable MPLS on the interface
    }
    rsvp {                                   # Required for dynamic MPLS only
        interface interface-name;
    }
}
```

- The general syntax descriptions provide no indication of the number of times you can specify a statement, option, or keyword. This information is provided in the text of the statement summary.

List of Technical Publications

Table 1 lists the software and hardware books for Juniper Networks routers and describes the contents of each book.

Table 1: Juniper Networks Technical Documentation

Book	Description
JUNOS Internet Software Configuration Guides	
Getting Started	Provides an overview of the JUNOS Internet software and describes how to install and upgrade the software. This manual also describes how to configure system management functions and how to configure the chassis, including user accounts, passwords, and redundancy.
Interfaces and Class of Service	Provides an overview of the interface and class-of-service functions of the JUNOS Internet software and describes how to configure the interfaces on the router.
MPLS Applications	Provides an overview of traffic engineering concepts and describes how to configure traffic engineering protocols.
Multicast	Provides an overview of multicast concepts and describes how to configure multicast routing protocols.
Network Management	Provides an overview of network management concepts and describes how to configure various network management features, such as SNMP, accounting options, and cflowd.
Policy Framework	Provides an overview of policy concepts and describes how to configure routing policy, firewall filters, and forwarding options.
Routing and Routing Protocols	Provides an overview of routing concepts and describes how to configure routing, routing instances, and unicast routing protocols.
VPNs	Provides an overview of Layer 2 and Layer 3 Virtual Private Networks (VPNs), describes how to configure VPNs, and provides configuration examples.
JUNOS Internet Software References	
Operational Mode Command Reference	Describes the JUNOS Internet software operational mode commands you use to monitor and troubleshoot Juniper Networks routers.
System Log Messages Reference	Describes how to access and interpret system log messages generated by JUNOS software modules and provides a reference page for each message.
JUNOScript API Documentation	
JUNOScript API Guide	Describes how to use the JUNOScript API to monitor and configure Juniper Networks routers.
JUNOScript API Reference	Provides a reference page for each tag in the JUNOScript API.

Book	Description
JUNOS Internet Software Comprehensive Index	
Comprehensive Index	Provides a complete index of all JUNOS Internet software books and the *JUNOScript API Guide*.
Hardware Documentation	
Hardware Guide	Describes how to install, maintain, and troubleshoot routers and router components. Each router platform (M5 and M10 routers, M20 router, M40 router, M40e router, M160 router, and T640 routing node) has its own hardware guide.
PIC Guide	Describes the router Physical Interface Cards (PICs). Each router platform has its own PIC guide.

Documentation Feedback

We are always interested in hearing from our customers. Please let us know what you like and do not like about the Juniper Networks documentation, and let us know of any suggestions you have for improving the documentation. Also, let us know if you find any mistakes in the documentation. Send your feedback to tech-doc@juniper.net.

How to Request Support

For technical support, contact Juniper Networks at support@juniper.net, or at 1-888-314-JTAC (within the United States) or 408-745-2121 (from outside the United States).

Part 1
Overview

- Command-Line Interface Overview on page 3
- Strategies for Monitoring and Troubleshooting the Router on page 31

Chapter 1

Command-Line Interface Overview

The *command-line interface (CLI)* is the interface to the software that you use whenever you access the router—whether from the console or through a remote network connection. The CLI, which automatically starts after the router finishes booting, provides commands that you use to perform various tasks, including configuring the JUNOS software and monitoring and troubleshooting the software, network connectivity, and the router hardware.

The CLI is a straightforward command interface. You type commands on a single line, and the commands are executed when you press the Enter key. The CLI provides command help and command completion, and it also provides Emacs-style keyboard sequences that allow you to move around on a command line and scroll through a buffer that contains recently executed commands.

The CLI is indicated by the presence of the > prompt, which is preceded by a string that defaults to the name of the user and the name of the router. For example:

```
user@host>
```

CLI Modes

The CLI has two modes: operational and configuration. In operational mode, you monitor and troubleshoot the software, network connectivity, and the router by entering CLI commands.

You configure the JUNOS software by entering a configuration mode and creating a hierarchy of configuration statements. For additional information about configuring the router, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

You monitor and troubleshoot the router by using the operational mode commands described in this manual.

This chapter discusses the operational CLI commands, including the following topics:

- CLI Command Hierarchy on page 4
- CLI Operational Commands on page 5
- Use the CLI on page 6
- Control the CLI Environment on page 16
- Set the Current Date and Time on page 18
- Display CLI Command History on page 19

- Monitor Who Uses the CLI on page 19
- Summary of CLI Operational Commands on page 20
- Summary of CLI Environment Commands on page 25

CLI Command Hierarchy

The CLI commands are organized in a hierarchical fashion, with commands that perform a similar function being grouped together under the same level. For example, all commands that display information about the system and the system software are grouped under the show command, and all commands that display information about the routing table are grouped under the show route command. Figure 1 illustrates a portion of the show command hierarchy.

Figure 1: CLI Command Hierarchy Example

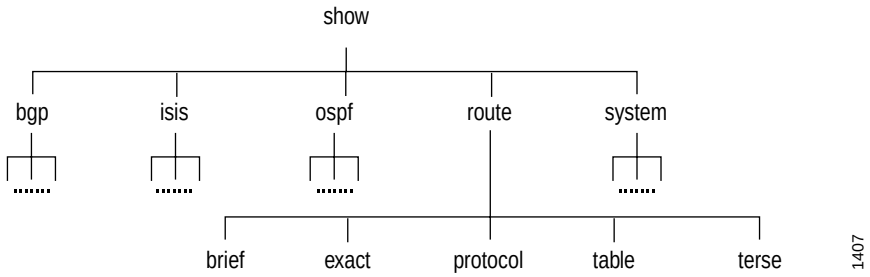

To execute a command, enter the full command name, starting at the top level of the hierarchy. For example, to display a brief view of the routes in the router table, use the command show route brief.

The hierarchical organization results in commands that have a regular syntax and provides several features that simplify use of the CLI:

- Consistent command names—Commands that provide the same type of function have the same name, regardless of the portion of the software they are operating on. For example, all show commands display software information and statistics, and all clear commands erase various types of system information.
- Lists and short descriptions of available commands—Information about available commands is provided at each level of the CLI command hierarchy. If you type a question mark (?) at any level, you see a list of the available commands along with a short description of each command. This means that if you already are familiar with the JUNOS software or with other routing software, you can use many of the CLI commands without consulting the documentation.
- Command completion—Command completion for command names (keywords) and for command options is available at each level of the hierarchy. If you type a partial command name followed immediately by a question mark (with no intervening space), you see a list of commands that match the partial name you typed.

CLI Operational Commands

When you first log into the router and the CLI starts, you are in CLI operational mode. At this level, there are several broad groups of CLI commands:

- Commands for controlling the CLI environment—The commands in the set hierarchy configure the CLI display screen. For information about these commands, see “Control the CLI Environment” on page 16.

- Monitoring and troubleshooting commands—The following commands let you display information and statistics about the software running on the router, such as routing table entries, and test network connectivity. Using these commands is discussed in various places in this manual.

  - clear—Clear protocol database information.

  - monitor—Perform real-time debugging of various software components, including the routing protocols and interfaces.

  - ping—Determine the reachability of a remote network host.

  - show—Display the configuration currently running on the router, and display information about the system, including interfaces, routing protocols, routing tables, routing policy filters, and the chassis.

  - test —Test the configuration and application of policy filters and AS path regular expressions.

  - traceroute—Trace the route to a remote network host.

- Commands for connecting to other network systems—The ssh command allows you to open secure shell connections and the telnet command allows you to open telnet sessions to other hosts on the network. For information about these commands, see “Monitor and Perform System Management Functions” on page 41.

- Commands for copying files—The file and copy commands allow you to copy files from one location on the router to another, from the router to a remote system, or from a remote system to the router. For information about these commands, see “Monitor and Perform System Management Functions” on page 41.

- Commands for restarting software processes—The commands in the restart hierarchy allow you to restart the various JUNOS software processes, including the routing protocol, interface, SNMP, and SNMP MIB II processes. For information about these commands, see the *JUNOS Internet Software Configuration Guide: Getting Started* and the *JUNOS Internet Software Configuration Guide: Network Management*.

- A command—request—for performing system-level operations, including stopping and rebooting the router and loading JUNOS software images. For information about this command, see “Monitor and Perform System Management Functions” on page 41.

- A command—start—to exit from the CLI and start a UNIX shell. For information about this command, see “Monitor and Perform System Management Functions” on page 41.

- A command—configure—for entering configuration mode, which provides a series of commands that let you configure the JUNOS software, including the routing protocols, interfaces, network management, and user access. For information about the CLI configuration commands, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

- A command—quit—for exiting from the CLI. For information about this command, see "Monitor and Perform System Management Functions" on page 41.

Use the CLI

This section describes how to use the *JUNOS* software CLI. It discusses the following topics:

- Get Help About Commands on page 6
- Have the CLI Complete Commands on page 8
- CLI Messages on page 9
- Move Around and Edit the Command Line on page 10

Get Help About Commands

The CLI provides context-sensitive help at every level of the command hierarchy. The help information tells you which commands are available at the current level in the hierarchy and provides a brief description of each.

To get help while in the CLI, type ?. You do not need to press Enter after typing the question mark. What you see on the screen depends on when and where you type the question mark:

- If you type the question mark at the command-line prompt, the CLI lists the available commands and options.

- If you type the question mark after entering the complete name of a command or command option, the CLI lists the available commands and options, then redisplays the command names and options that you typed.

- If you type the question mark in the middle of a command name, the CLI lists possible command completions that match the letters you have entered so far, then redisplays the letters that you typed.

Examples: Get Help with Commands

In this example, all available commands are listed at the top level of the command hierarchy:

```
user@host> ?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
  file           Perform file operations
  help           Provide help information
  monitor        Real-time debugging
  mtrace         Trace multicast path from a source to a receiver
  ping           Ping a remote target
  quit           Exit the management session
  request        Make system-level requests
  restart        Restart a software process
  set            Set CLI properties, date, time, craft display text
  show           Show information about the system
  ssh            Open a secure shell to another host
  start          Start a software process
  telnet         Telnet to another host
  test           Diagnostic debugging commands
  traceroute     Trace the route to a remote host
user@host>
```

List all commands that start with the letter c:

```
user@host> c?
Possible completions:
  clear          Clear information in the system
  configure      Manipulate software configuration information
user@host> c
```

List all available clear commands:

```
user@host> clear ?
Possible completions:
  arp            Clear address-resolution information
  bgp            Clear BGP information
  chassis        Clear chassis information
  firewall       Clear firewall counters
  igmp           Clear IGMP information
  interfaces     Clear interface information
  isis           Clear IS-IS information
  ldp            Clear LDP information
  log            Clear contents of a log file
  mpls           Clear MPLS information
  msdp           Clear MSDP information
  multicast      Clear Multicast information
  ospf           Clear OSPF information
  pim            Clear PIM information
  rip            Clear RIP information
  route          Clear routing table information
  rsvp           Clear RSVP information
  snmp           Clear SNMP information
  system         Clear system status
  vrrp           Clear VRRP statistics information
user@host> clear
```

Have the CLI Complete Commands

You do not always have to remember or type the full command or option name for the CLI to recognize it. To display all possible command or option completions, type the partial command followed immediately by a question mark.

To complete a command or option that you have partially typed, press Tab or the Space Bar. If the partially typed letters begin a string that uniquely identifies a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command, and the screen displays the possible completions.

Command completion also applies to others strings, such as filenames and usernames. To display all possible values, you can type a partial string followed immediately by a question mark. However, to complete these strings, press Tab; pressing the Space Bar does not work.

Examples: Use CLI Command Completion

Issue the show interface command:

```
user@host> sh<Space>ow i<Space>
'i' is ambiguous.
Possible completions:
   igmp        Show information about IGMP
   interface   Show interface information
   isis        Show information about IS-IS
user@host> show in<Space>terfaces <Enter>
Physical interface: at-0/1/0, Enabled, Physical link is Up
  Interface index: 11, SNMP ifIndex: 65
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode
  Speed: OC12, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : 0x01
...
user@host>
```

Display a list of all log files whose name starts with the string "messages," and then display the contents of one of the files:

```
user@myhost> show log mes?
Possible completions:
  <filename>        Log file to display
  messages          Size: 1417052, Last changed: Mar  3 00:33
  messages.0.gz     Size: 145575, Last changed: Mar  3 00:00
  messages.1.gz     Size: 134253, Last changed: Mar  2 23:00
  messages.10.gz    Size: 137022, Last changed: Mar  2 14:00
  messages.2.gr     Size: 137112, Last changed: Mar  2 22:00
  messages.3.gz     Size: 121633, Last changed: Mar  2 21:00
  messages.4.gz     Size: 135715, Last changed: Mar  2 20:00
  messages.5.gz     Size: 137504, Last changed: Mar  2 19:00
  messages.6.gz     Size: 134591, Last changed: Mar  2 18:00
  messages.7.gz     Size: 132670, Last changed: Mar  2 17:00
  messages.8.gz     Size: 136596, Last changed: Mar  2 16:00
  messages.9.gz     Size: 136210, Last changed: Mar  2 15:00
user@myhost> show log mes<Tab>sages.4<Tab>.gz<Enter>
Jan 15 21:00:00 myhost newsyslog[1381]: logfile turned over
...
```

CLI Messages

You see messages when you enter and exit from configuration mode, when you commit a configuration, and when you type a string or value that is not valid.

When you commit a configuration, the JUNOS software checks the configuration you are committing. If there are no problems, a message indicates that the configuration was accepted. If there are problems, a message indicates where the errors are.

In the top-level CLI commands and in configuration mode, if you type an invalid string—for example, the name of a command or statement that does not exist—you see the message syntax error or unknown command. A caret (^) indicates where the error is. Examples:

```
user@host> clear route
                       ^
syntax error, expecting <command>.

[edit]
user@host# telnet
           ^
unknown command.
```

When the number of choices is limited, the commands you can enter to correct the syntax error might be displayed. For example:

```
[edit]
user@host# load myconfig-file<Enter>
                ^
syntax error, expecting ‘merge’, ‘override’, or ‘replace’.
```

Move Around and Edit the Command Line

In the CLI, you can use keyboard sequences to move around on a command line and edit the command line. You can also use keyboard sequences to scroll through a list of recently executed commands. Table 2 lists the CLI keyboard sequences. They are the same as those used in Emacs.

Table 2: CLI Keyboard Sequences

Category	Action	Keyboard Sequence
Move the Cursor	Move the cursor back one character.	Ctrl-b
	Move the cursor back one word.	Esc-b or Alt-b
	Move the cursor forward one character.	Ctrl-f
	Move the cursor forward one word.	Esc-f or Alt-f
	Move the cursor to the beginning of the command line.	Ctrl-a
	Move the cursor to the end of the command line.	Ctrl-e
Delete Characters	Delete the character before the cursor.	Ctrl-h, Delete, or Backspace
	Delete the character at the cursor.	Ctrl-d
	Delete all characters from the cursor to the end of the command line.	Ctrl-k
	Delete all characters on the command line.	Ctrl-u or Ctrl-x
	Delete the word before the cursor.	Ctrl-w, Esc-Backspace, or Alt-Backspace
	Delete the word after the cursor.	Esc-d or Alt-d
Insert Recently Deleted Text	Insert the most recently deleted text at the cursor.	Ctrl-y
Redraw the Screen	Redraw the current line.	Ctrl-l
Display Previous Command Lines	Scroll backward through the list of recently executed commands.	Ctrl-p
	Scroll forward through the list of recently executed commands.	Ctrl-n
	Search the CLI history in reverse order for lines matching the search string.	Ctrl-r
	Search the CLI history by typing some text at the prompt, followed by the keyboard sequence. The CLI attempts to expand the text into the most recent word in the history for which the text is a prefix.	Esc-/
Repeat Keyboard Sequences	Specify the number of times to execute a keyboard sequence. *number* can be from 1 through 9.	Esc-*number sequence* or Alt-*number sequence*

How Output Appears on the Screen

When you issue commands in operational mode CLI, the output appears on the screen. If the output is longer than the screen, you can display it one screen at a time using a UNIX more-type interface.

You can also redirect the output through a UNIX-type pipe, either to apply simple filters to the output or to redirect the output to a file.

Display Output One Screen at a Time

If the output is longer than the screen, the ---More--- prompt indicates that more output is available. Table 3 describes how you can scroll at the ---More--- prompt. As soon as the CLI calculates how long the output is (usually by the second screen), the percentage of the command output already displayed appears next to the prompt.

Table 3: ---More--- Prompt Keyboard Sequences

Objective	Result	Keyboard Sequence
Get Help	Display help information about the keyboard sequences you can display at the ---More--- prompt.	h
Scroll Down	Scroll down one line.	Enter, Return, k, Ctrl-m, Ctrl-n, or down arrow
	Scroll down one-half screen.	Tab, d, Ctrl-d, or Ctrl-x
	Scroll down one whole screen.	Space or Ctrl-f
	Scroll down to the bottom of the output.	Ctrl-e or G
	Display the output all at once instead of one screen at a time. (Same as specifying the | no-more command after a pipe.)	N
Scroll Up	Display the previous line of output.	j, Ctrl-h, Ctrl-p, or up arrow
	Scroll up one-half screen.	u or Ctrl-u
	Scroll up one whole screen.	b or Ctrl-b
	Scroll up to the top of the output.	Ctrl-a or g
Search	Search forward for a string.	/*string*
	Search backward for a string.	?*string*
	Repeat the previous search for a string.	n
	Search for a text string. You are prompted for the string to match. (Same as specifying the | match *string* command.)	m or M
	Search, ignoring a text string. You are prompted for the string to not match. (Same as specifying the | except *string* command.)	e or E
Interrupt or End Output, Redraw the Output, and Save the Output to a File	Interrupt the display of output.	Ctrl-C, q, Q, or Ctrl-k
	Do not redisplay the CLI prompt immediately after displaying the output, but remain at the ---More--- prompt. (Same as specifying the | hold command.)	H
	Clear any match conditions and display the complete output.	c or C
	Redraw the output on the screen.	Ctrl-l
	Save the command output to a file. You are prompted for a filename. (Same as specifying the | save *filename* command.)	s or S

Filter Command Output

You can redirect the output of operational mode commands into a file or into filters. When you display help about these commands, one of the options listed is |, which is a pipe to redirect the output. For example:

```
user@host> show configuration ?
Possible completions:
  <[Enter]>  Execute this command
  |          Pipe through a command

user@host> show configuration | ?
Possible completions:
  count      Count occurrences
  except     Show only text that does not match a pattern
  find       Search for the first occurrence of a pattern
  hold       Hold text without exiting the ---More--- prompt
  match      Show only text that matches a pattern
  no-more    Don't paginate output
  resolve    Resolve IP addresses
  save       Save output text to a file
  trim       Trim specified number of columns from the start line
```

Place Command Output in a File

When the output is very long, when you need to store or analyze the output, or when you need to send the output in e-mail, you can redirect the output to a file. Doing this is useful when the output scrolls off the screen, making it difficult to cut the output from a window and paste it into another.

To save the output to a file, specify the save command after the pipe:

```
user@host> command | save filename
```

By default, the file is placed in your home directory on the router. For information about how you can specify the filename, see “How to Specify Filenames and URLs” on page 43.

This example stores the output of the request support information command in a file:

```
user@host> request support information | save filename
Wrote 1143 lines of output to ‘filename’
user@host>
```

Search for a String in the Output

You can filter the output to search for a text matching a regular expression. You can match a regular expression, match everything except a regular expression, or find the first occurrence of text matching a regular expression. All searches are not case-sensitive.

To match a regular expression, specify the match command after the pipe:

user@host> *command* **| match** *regular-expression*

To ignore text that matches a regular expression, specify the except command after the pipe:

user@host> *command* **| except** *regular-expression*

If the *regular-expression* contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

You use extended regular expressions to specify what text in the output to match. Command regular expressions implement the extended (modern) regular expressions as defined in POSIX 1003.2. Table 4 lists common regular expression operators.

Table 4: Common Regular Expression Operators

Operator	Match...
|	One of the two terms on either side of the pipe.
^	At the beginning of an expression, used to denote where the command begins, where there might be some ambiguity.
$	Character at the end of a command. Used to denote a command that must be matched exactly up to that point. For example, allow-commands "show interfaces $" means that the user cannot issue show interfaces detail or show interfaces extensive.
[ ]	Range of letters or digits. To separate the start and end of a range, use a hyphen (-).
( )	A group of commands, indicating an expression to be evaluated and the result is then evaluated as part of the overall expression.

For example, if a command produces the following output:

```
one two
two two
three two one
four
```

The match two command displays:

```
one two
two two
three two one
```

The except one command displays:

```
two two
four
```

List all the ATM interfaces in the configuration:

```
user@host> show configuration | match at-
   at-2/1/0 {
   at-2/1/1 {
   at-2/2/0 {
   at-5/2/0 {
   at-5/3/0 {
```

Display a skeleton of your router configuration:

```
[edit]
user@host # show | match {
system {
   root-authentication {
   name-server {
   login {
      class superuser {
      user junipero {
         authentication {
   services {
   syslog {
      file messages {
   processes {
chassis {
   alarm {
      sonet {
   images {
      scb {
      fpc {
interfaces {
   at-2/1/1 {
      atm-options {
      unit 0 {
   at-2/2/0 {
   ...
snmp {
   community public {
      clients {
routing-options {
   static {
      route 0.0.0.0/0 {
      route 192.168.0.0/16 {
      route 208.197.169.0/24 {
protocols {
   rsvp {
      interface so-5/1/0 {
   mpls {
      interface so-5/1/0 {
   bgp {
      group internal {
   ospf {
      area 0.0.0.0 {
         interface so-5/1/0 {
```

List all users who are logged into the router except for the user “root”:

```
user@host> show system users | except root
 8:28PM  up 1 day, 13:59, 2 users, load averages: 0.01, 0.01, 0.00
USER     TTY FROM              LOGIN@  IDLE WHAT
sheep    p0  baa.juniper.net   7:25PM     - cli
```

Save the configuration, except for encrypted passwords, to a file:

```
user@host> show configuration | except SECRET-DATA | save my.output.file
```

Display the output, starting not at the beginning but rather at the first occurrence of text matching a regular expression, using the find command after the pipe:

```
user@host> command | find regular-expression
```

If the regular expression contains spaces, operators, or wildcard characters, enclose the expression in quotation marks.

List the routes in the routing table starting at 208.197.169.0:

```
user@host> show route | find 208.197.169.0
208.197.169.0/24   *[Static/5] 1d 13:22:11
                    > to 192.168.4.254 via so-3/0/0.0
224.0.0.5/32       *[OSPF/10] 1d 13:22:12, metric 1

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.4015.00/160
                   *[Direct/0] 1d 13:22:12
                    > via lo0.0
```

Count the Number of Lines in the Output

To count the number of lines in the output, specify the count command after the pipe:

```
user@host> command | count
```

For example:

```
user@host> show configuration | count
Count: 269 lines
user@host> show route | count
Count: 67 lines
```

Display All Output at Once

To display the output all at once instead of one screen at a time, specify the no-more command after the pipe. This command is equivalent to the set cli screen-length 0 command, but affects the output of the one command only.

```
user@host> command | no-more
```

Retain the Output after the Last Screen

When you view output one screen at a time, you typically return to the CLI prompt after viewing the last screen.

To not return immediately, use the hold command after the pipe. This feature is useful, for example, when you want to scroll or search through the output.

```
user@host> command | hold
```

Filter Command Output Multiple Times

For the output of a single command, you can redirect the output one or more times. For example:

```
user@host> command | match regular-expression | except regular-expression |
match other-regular-expression | find regular-expression | hold
```

Control the CLI Environment

To configure the CLI environment, use the CLI operational set command:

```
user@host> set cli ?
Possible completions:
  complete-on-space    Toggle word completion on space
  idle-timeout         Set the cli maximum idle time
  prompt               Set the cli command prompt string
  restart-on-upgrade   Set cli to prompt for restart after a software
                        upgrade
  screen-length        Set number of lines on screen
  screen-width         Set number of characters on a line
  terminal             Set terminal type
```

When you log in to the router using ssh, or log in from the console when its terminal type has already been configured, your terminal's type, screen length, and screen width are already set, so you do not need to change them from the CLI.

Set the Terminal Type

To set the terminal type, use the set cli terminal command:

```
user@host> set cli terminal terminal-type
```

The *terminal-type* can be one of the following: ansi, vt100, small-xterm, xterm.

Set the Screen Length

The default CLI screen length is 24 lines. To change the length, use the set cli screen-length command:

```
user@host> set cli screen-length length
```

Setting the screen length to 0 lines disables the display of output one screen at a time. Disabling this UNIX more-type interface can be useful when you are issuing CLI commands from scripts.

Set the Screen Width

The default CLI screen width is 80 columns. To change the length, use the set cli screen-width command:

```
user@host> set cli screen-width width
```

Set the CLI Prompt

The default CLI prompt is *user@host*>. To change this, use the set cli prompt command. If the prompt string contains spaces, enclose the string in quotation marks.

```
user@host> set cli prompt string
```

Idle Timeout

By default, an individual CLI session never times out, unless the idle-timeout statement has been included in the user's login class configuration. To set the maximum time an individual session can be idle before the user is logged off the router, use the set cli idle-timeout command:

```
user@host> set cli idle-timeout timeout
```

timeout can be 0 to 100,000 minutes. Setting *timeout* to 0 disables the timeout.

Set CLI to Prompt after a Software Upgrade

By default, the CLI prompts you to restart after a software upgrade. To disable the prompt for an individual session, use the set cli restart-on-upgrade off command:

```
user@host> set cli restart-on-upgrade off
```

To re-enable the prompt, use the set cli restart-on-upgrade on command:

```
user@host> set cli restart-on-upgrade on
```

Set Command Completion

By default, you can type a space or tab to have the CLI complete a command. To have the CLI allow only a tab to complete a command, use the set cli complete-on-space command:

```
user@host> set cli complete-on-space off
Disabling complete-on-space
user@host>
```

To re-enable the use of both space and tab characters for command completion, use the set cli complete-on-space on command:

```
user@host> set cli complete-on-space on
Enabling complete-on-space
user@host>
```

Display CLI Settings

To display the current CLI settings, use the show cli command:

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 24
CLI screen-width set to 80
CLI terminal is 'ansi'
```

Example: Control the CLI Environment

The following example shows how to change the default CLI environment:

```
user@host> set cli screen-length 66
Screen length set to 66
user@host> set cli screen-width 40
Screen width set to 40
user@host> set cli prompt “router1-san-jose > “
router1-san-jose > show cli
CLI screen length set to 66
CLI screen width set to 40
router1-san-jose >
```

Set the Current Date and Time

To set the current date and time on the router, use the set date command:

```
user@host> set date YYYYMMDDhhmm.ss
```

YYYY is the four-digit year, *MM* is the two-digit month, *DD* is the two-digit date, *hh* is the two-digit hour, *mm* is the two-digit minute, and *ss* is the two-digit second. At a minimum, you must specify the two-digit minute. All other parts of the date and time are optional.

For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Set the Current Date and Time from NTP Servers

To synchronize the current date and time on the router to NTP servers, use the set date ntp command:

```
user@host> set date ntp <ntp-server>
```

You do not need to reboot the router when you use the set date ntp command. If you do not specify any NTP servers, the JUNOS software will use the ones configured at the [edit system ntp server] hierarchy level.

For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Display CLI Command History

You can display a list of recent commands that you issued. To display the command history, use the show cli history command:

```
user@host> show cli history
  03-03 01:00:50 -- show cli history
  03-03 01:01:12 -- show interfaces terse
  03-03 01:01:22 -- show interfaces lo0
  03-03 01:01:44 -- show bgp next-hop-database
  03-03 01:01:51 -- show cli history
```

By default, this command displays the last 100 commands issued in the CLI. If you specify a number with the command, it displays that number of recent commands. For example:

```
user@host> show cli history 3
  01:01:44 -- show bgp next-hop-database
  01:01:51 -- show cli history
  01:02:51 -- show cli history 3
```

Monitor Who Uses the CLI

Depending upon how you configure the JUNOS software, multiple users can log into the router, use the CLI, and configure or modify the software configuration.

The JUNOS software provides a general syslog-like mechanism to log system operations, such as when users log into the router and when they issue CLI commands. To configure system logging, include the syslog statement in the configuration, as described in the *JUNOS Internet Software Configuration Guide: Getting Started*.

If another user is in configuration mode when you enter configuration mode, a notification message is displayed that indicates who the user is and what portion of the configuration they are viewing or editing:

```
user@host> configure
Entering configuration mode
Current configuration users:
  root terminal p3 (pid 1088) on since 1999-05-13 01:03:27 EDT
    [edit interfaces fxp0 unit 0 family inet]
The configuration has been changed but not committed
```

Summary of CLI Operational Commands

The following sections explain each of the top-level CLI commands. The commands are organized alphabetically.

clear

Syntax clear (arp | bgp | chassis | firewall | igmp | ilmi | interfaces | isis | ldp | log | mpls | msdp | multicast | ospf | pim | rip | route | rsvp | snmp | system | vrrp)

Description Delete or zero information from various router and protocol databases.

Usage Guidelines The various clear commands are discussed in other chapters. To locate these commands, refer to the table of contents or the index.

Required Privilege Level clear

configure

Syntax configure

Description Enter configuration mode.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level configure

file

Syntax file (copy | delete | list | rename | show)

Description Copy files to and from the router.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level maintenance

help

Syntax help (reference | topic)

Description Display help information about available operational commands and configuration statements.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level view

monitor

Syntax monitor (start | stop | list) <interface > <traffic>

Description Monitor a log file or interface statistics or traffic in real time.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41, “Real-Time Router Interface Monitoring and Troubleshooting” on page 445, and monitor traffic on page 452.

Required Privilege Level Depends on the specific command.

mtrace

Syntax mtrace (from-source | to-gateway | monitor)

Description Display trace information about a multicast path from a source to a receiver.

Usage Guidelines See “IP Multicast Monitoring and Troubleshooting” on page 639.

Required privilege level view

ping

Syntax ping

Description Check the reachability of network hosts.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41.

Required Privilege Level network

quit

Syntax quit

Description Exit from the CLI to a UNIX shell.

Required Privilege Level shell and maintenance

See Also start on page 24

| (pipe)

Syntax | (compare | count | display <detail | inheritance | xml> | except *pattern* | find *pattern* | hold | match *pattern* | no-more | resolve <full-names> | save *filename* | trim *columns*)

Description Filter the output of an operational mode or a configuration mode command.

Options compare (filename | rollback n)—(Configuration mode only, with the show command only) Compare configuration changes with another configuration file.

count—Display the number of lines in the output.

display—Display additional information about the contents of the configuration.

- detail—(Configuration mode only) Display configuration data detail.
- inheritance—(Configuration mode only) Display inherited configuration data and source group.
- xml—(Operational mode only) Display XML content of the command.

except *pattern*—Ignore text matching a regular expression when searching the output. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.

find *pattern*—Display the output starting at the first occurrence of text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.

hold—Hold text without exiting the --More-- prompt.

match *pattern*—Search for text matching a regular expression. If the regular expression contains spaces, operators, or wildcard characters, enclose it in quotation marks.

no-more—Display output all at once rather than one screen at a time.

resolve—Convert IP addresses into DNS names. Truncates to fit original size unless full-names specified. To prevent the names from being truncated, use the full-name option.

save *filename*—Save the output to a file or URL. For information about specifying the filename, see "Search for a String in the Output" on page 13.

trim *columns*—Trim specified number of columns from the start line.

Usage Guidelines See "Filter Command Output" on page 12.

request

Syntax request system (reboot | halt | software | snapshot)

Description Stop or reboot the router, load software packages, and back up the router's file systems.

Usage Guidelines See "Monitor and Perform System Management Functions" on page 41.

Required Privilege Level maintenance

restart

Syntax restart (fpc | interface-control | mib-process | routing | sampling | sfm | snmp | soft)

Description Restart router software processes.

Usage Guidelines See "Monitor and Perform System Management Functions" on page 41.

Required Privilege Level reset

set

Syntax set (chassis | cli | date | date ntp)

Description Configure CLI properties and the router's date and time.

Usage Guidelines See "Control the CLI Environment" on page 16, "Set the Current Date and Time" on page 18, and "Set the Current Date and Time from NTP Servers" on page 18.

Required Privilege Level view

show

Syntax show (aps | arp | as-path | bgp | chassis | cli | configuration | connections | dvmrp | firewall | host | igmp | interfaces | isis | ldp | log | mpls | msdp | multicast | ntp | ospf | pfe | pim | policy | rip | route | rsvp | sap | snmp | system | task | ted | version | vrrp)

Description Show information about all aspects of the software, including interfaces and the routing protocols.

Usage Guidelines The various show commands are discussed on other chapters. To locate these commands, refer to the table of contents or the index.

Required Privilege Level Depends on the specific command.

ssh

Syntax ssh

Description Open a secure shell to another host.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41.

Required Privilege Level network

start

Syntax start shell

Description Start a UNIX shell on the router.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41.

Required Privilege Level shell and maintenance

telnet

Syntax telnet

Description Establish a telnet session to another host.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41.

Required Privilege Level network

test

Syntax test (configuration | interface | msdp | policy)

Description Run various diagnostic debugging commands.

Usage Guidelines The various test commands are discussed in other chapters. To locate these commands, refer to the table of contents or the index.

Required Privilege Level Depends on the specific command.

traceroute

Syntax traceroute

Description Trace the route to a remote host.

Usage Guidelines See “Monitor and Perform System Management Functions” on page 41.

Required Privilege Level network

Summary of CLI Environment Commands

The following sections explain each of the CLI environment commands. The commands are organized alphabetically.

set cli complete-on-space

Syntax set cli complete-on-space (on | off);

Description Configure the keys to use for command completion.

Default When you type a space or tab, the CLI performs command completion.

Options on—Allow both a space and a tab to be used for command completion.

off—Allow only a tab to be used for command completion.

Usage Guidelines See "Set Command Completion" on page 17.

Required Privilege Level view

Sample Output

```
user@host> set cli com<Space>
user@host> set cli complete-on-space off
user@host> set cli com<Tab>
user@host> set cli complete-on-space on
user@host>
```

set cli idle-timeout

Syntax set cli idle-timeout *timeout*

Description Set the maximum time that an individual session can be idle before the user is logged off the router. The session times out after remaining at the CLI operational mode prompt for the specified time. The session can time out while monitoring log files.

Default If you do not issue this command, and the user's login class does not specify this value, the user is never forced off the system after extended idle times.

Options *timeout*—Maximum idle time, in minutes

Range: 0 through 100,000 minutes. Setting it to 0 disables the timeout.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level view

set cli prompt

Syntax set cli prompt *string*

Description Set the prompt to display within the CLI.

Default *user@host>*

Options *string*—CLI prompt. To include spaces in the prompt, enclose the string in quotation marks.

Usage Guidelines See "Set the CLI Prompt" on page 17.

Required Privilege Level view

Sample Output

```
user@host> set cli prompt "cli% "
cli%
```

set cli restart-on-upgrade

Syntax set cli restart-on-upgrade (off | on)

Description For an individual session, set the CLI to prompt you to restart the router after upgrading the software.

Default The CLI prompts you to restart, unless the screen length has been set to 0.

Options off—Disables the prompt.

on—Enables the prompt.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level view

set cli screen-length

Syntax set cli screen-length *length*

Description Set the number of lines of text that the screen can display.

Options *length*—Number of lines on the screen.
Range: 0 through 100,000
Default: 24 lines

Usage Guidelines See "Set the Screen Length" on page 16.

Required Privilege Level view

Sample Output

```
user@host> set cli screen-length 66
Screen length is set to 66
user@host>
```

set cli screen-width

Syntax set cli screen-width *width*

Description Set the number of characters that the screen can display on a single line.

Options *width*—Number of columns on the screen.
Range: 0 through 100,000
Default: 80 columns

Usage Guidelines See "Set the Screen Width" on page 16.

Required Privilege Level view

Sample Output

```
user@host> set cli screen-width 40
Screen width set to 40
user@host>
```

set cli terminal

Syntax set cli terminal *terminal-type*

Description Set the terminal type.

Options *terminal-type*—Type of terminal that is connected to the port.
Values: ansi, vt100, small-xterm, xterm
Default: The terminal type is unknown, and the user is prompted for the terminal type.

Usage Guidelines See "Set the Terminal Type" on page 16.

Required Privilege Level view

set date

Syntax set date *YYYYMMDDhhmm.ss*

Description Set the current date and time on the router.

Options *YYYYMMDDhhmm.ss*—Date and time to set. *YYYY* is the four-digit year, *MM* is the two-digit month, *DD* is the two-digit date, *hh* is the two-digit hour, *mm* is the two-digit minute, and *ss* is the two-digit second. At a minimum, you must specify the two-digit minute. All other parts of the date and time are optional.

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level view

set date ntp

Syntax set date ntp <*ntp-server*>

Description Use an NTP server to synchronize the current date and time setting on the router.

You do not need to reboot the router when you use the set date ntp command.

Options none—Uses the system NTP server list.

ntp-server—IP address of one or more NTP servers to query. When querying more than one server, use the format "*ip-address ip-address*" (for example, "1.2.3.4 192.17.22.9").

Required Privilege Level view

Usage Guidelines See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Sample Output: set date ntp (to query one server)

```
user@host> set date ntp 192.17.12.9
14 Sep 22:00:50 ntpdate[20603]: step time server 192.17.12.9 offset 0.000461 sec
```

Sample Output: set date ntp (to query two servers)

```
user@host> set date ntp "200.49.40.1 129.127.28.4"
10 Feb 13:50:21 ntpdate[794]: step time server 129.127.28.4 offset 0.000163 sec
```

show cli

Syntax show cli

Description Display information about how the CLI environment is configured.

Required Privilege Level view

Usage Guidelines See "Control the CLI Environment" on page 16.

Sample Output

```
user@host> show cli
CLI screen length set to 60
CLI screen width set to 80
CLI complete-on-space set to on
user@host>
```

show cli history

Syntax show cli history <*count*>

Description List recent commands that you issued in the CLI and the time they were issued.

If you issue the run show cli history command from configuration mode, the command lists the most recent configuration mode commands that you issued and the time they were issued.

Options *count*—(Optional) Number of recent commands to display.
Range: 0 through 65,535
Default: 100 (20 if your screen length is set to 0)

Usage Guidelines See "Display CLI Command History" on page 19.

Required Privilege Level view

Sample Output

```
user@host> show cli history
  12:33:39 -- configure
  12:42:52 -- show cli history
  12:43:02 -- show interfaces terse
  12:43:14 -- show interfaces lo0
  12:43:20 -- show bgp
  12:43:28 -- show bgp next-hop-database
  12:43:32 -- show cli history
user@host> configure
...
[edit]
user@host# run show cli history
  12:40:08 -- show
  12:40:17 -- edit protocols
  12:40:27 -- set isis
  12:40:29 -- edit isis
  12:40:40 -- run show cli history
[edit protocols isis]
user@host#
```

Chapter 2

Strategies for Monitoring and Troubleshooting the Router

This chapter describes the standard methods and most common tools used in troubleshooting the router. It discusses troubleshooting basics, including how, when, and why you use commands to monitor and troubleshoot the router and network. This chapter also includes a list of the commands most commonly used by Network Operation Centers (NOCs) for monitoring and troubleshooting.

This chapter discusses the following topics:

- Basic Approaches to Troubleshooting on page 31
- Tools for Troubleshooting on page 36

Basic Approaches to Troubleshooting

This section discusses the following aspects of troubleshooting:

- Troubleshooting Process on page 31
- Identify the Symptoms on page 32
- Isolate the Cause on page 32
- Take Corrective Action on page 33
- Evaluate the Solution on page 33

For a troubleshooting example, see “Example: Strategy for Isolating a Broken Network Connection” on page 34.

Troubleshooting Process

Troubleshooting can be simplified by following standard procedures, as illustrated by Figure 2. Standard troubleshooting procedures include the following steps:

1. Identify symptoms. A symptom can be defined as any unwanted results or behavior. A problem or failure might exhibit one or more symptoms.
2. Isolate the cause of the symptom.
3. Take action to correct the problem.

4. Evaluate the system to see if the original problem is solved and to verify that new problems have not been introduced by the changes you made.

5. At this point, either you have solved the problem, or you must return to Steps 1, 2, or 3 and identify the symptoms more clearly, isolate additional possible causes, or take additional action to correct the problem.

   During the troubleshooting process, you might isolate causes that require additional troubleshooting before you can continue with the standard process.

Figure 2: Troubleshooting Process

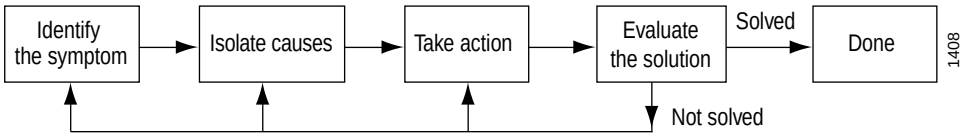

Identify the Symptoms

Identifying symptoms requires careful observation. The best preparation for troubleshooting is knowing your network thoroughly before a problem occurs so that you have a baseline state from which to work. If you understand how the network functions under normal conditions, it is easier to distinguish between normal and abnormal activity.

Sometimes a problem is related to another condition that must be solved first. When identifying symptoms, record as many parameters as you can regarding the offending state. The more information you have, the easier it is to isolate the cause. If you find a set of symptoms, try to decide what they have in common. It is likely that they are related, and noticing as many symptoms as you can provides you with more information as you proceed.

It is also useful to record what changes have taken place since the system was last functioning correctly. Changes in activity are likely to be related to changes in configuration.

Isolate the Cause

A particular symptom can be the result of one or more causes. Successful troubleshooting requires narrowing the focus to find each individual cause for unwanted behavior. While you might find a solution by just trying a variety of actions, you reach the intended solution more quickly if you systematically approach the problem.

There are several useful methods for isolating a problem:

- Retrace your steps—Try to return to a state that existed before the problem appeared. When the network is in a known state, take small steps forward, watching carefully for the recurrence of symptoms.

- Divide the problem into its smallest unit—Cut the problem in half and test each half. If only one half continues to have the problem, cut it in half again or compare it to the valid half to see how it is different. You might find the solution in the difference.

- Identify which functions are working correctly—Do not waste time investigating functions that are not broken.

- Keep careful records of changes and effects—Ask questions and document changes as you work on a system.
- Notice how various symptoms might be related—If you are finding unexpected or unwanted results in more than one area, try to discover what those areas have in common and what variables would affect them. You probably will find the source for the problem in the common areas.
- Imagine what type of errors or failures could lead to the particular symptom—Test for the errors or failures to see if they are actually occurring.
- Do not try to solve multiple unrelated problems simultaneously—If multiple symptoms occur that do not appear to be related, select one symptom or set of symptoms and focus on it. However, do not completely ignore the other symptoms, because you might discover that they are related after all.

Several useful tools exist for isolating the cause of a problem, including network analyzer traces, core dumps, serial line traces, stack dumps, and the output from various show commands in CLI. For information about the show commands, see the chapters in this manual that describe the JUNOS monitoring commands.

Take Corrective Action

The action required depends on the type of problem you have isolated. As you troubleshoot, keep in mind the following principles:

- Document each step you take.
- Use the various CLI show commands to verify which behaviors change with each action you take.
- When you are considering several possible actions, you can choose to test the easiest first, thereby eliminating possibilities quickly, or you can choose the action that appears most likely to solve the problem, even if it is more time-consuming or difficult to perform.

Evaluate the Solution

Carefully test the solution to ensure that it does not introduce new symptoms. If new symptoms occur, start the troubleshooting process again, carefully documenting the changes you make in the process.

Example: Strategy for Isolating a Broken Network Connection

To illustrate the troubleshooting process, we examine a problem that appears to include a broken network connection. By applying the strategy listed below and shown in Figure 3, you can usually isolate the failed node:

- Identify symptom—Failure to reach remote host.
- Isolate causes—Several possible causes are identified, including:
  - Local router is misconfigured.
  - Intermediate router is misconfigured.
  - Remote router is misconfigured.
  - No path to the remote router in the local routing table.
- Take action—Appropriate actions are taken for each possible cause, including:
  - Check the local router's configuration and edit if appropriate.
  - Troubleshoot intermediate router.
  - Check the remote host configuration and edit if appropriate.
  - Troubleshoot routing protocols.
  - Identify additional possible causes.
- Evaluate solution—If the problem is solved, you are done. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In Figure 3, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you had reason to believe the problem was related to a known issue, such as a recent change in configuration.

Often, troubleshooting one symptom will uncover other symptoms. Figure 3 shows two possible causes that might involve additional troubleshooting.

Figure 3: Example: Troubleshooting Loss of Connectivity

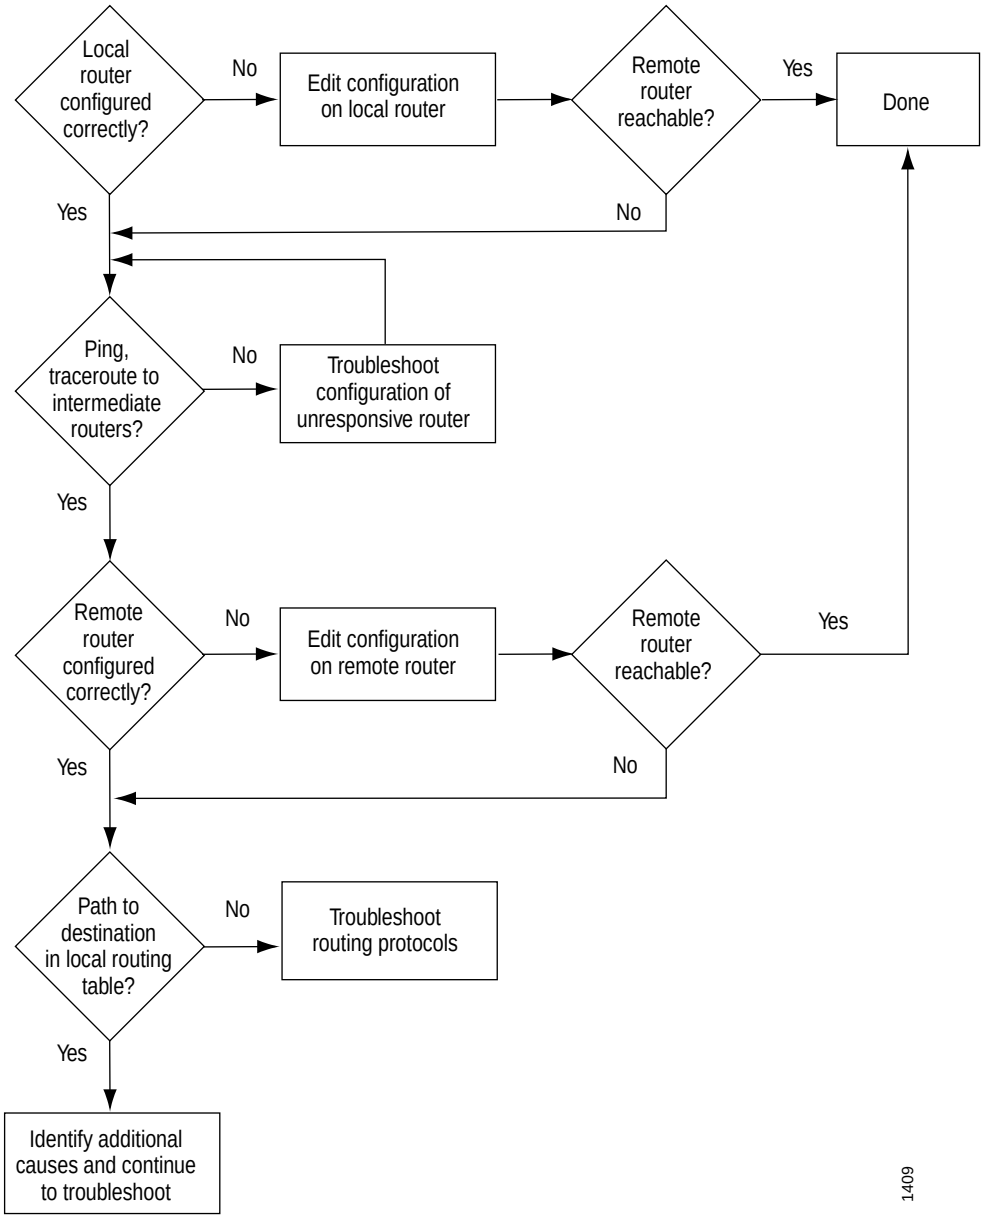

Tools for Troubleshooting

You can use the following tools when troubleshooting your network:

- Commonly Used Operational Mode Commands on page 36
- Craft Interface on page 37

Commonly Used Operational Mode Commands

Table 5 lists common operational mode commands that you can use to monitor and troubleshoot the router, network, and traffic.

Table 5: Commonly Used Operational Mode Commands

Task	Task or Information to Monitor	Command
Software Version	Versions of software running on the router.	show version on page 109
Log Files	Contents of the log files.	monitor on page 47
	Log files and their contents and recent user logins.	show log on page 69
Remote System Reachability	Host reachability and network connectivity.	ping on page 48
	Route to a network system.	traceroute on page 114
Configuration	Current running system configuration.	show configuration on page 67
Manipulate Files	List of files and directories on the router.	file list on page 46
	Contents of a file.	file show on page 46
Static Interface Information	Detailed information about interfaces.	show interfaces on page 194
	Summary information about interfaces.	show interfaces terse on page 206
Chassis	Information currently on craft display.	show chassis craft-interface on page 136
	Router environment information.	show chassis environment on page 138
	Hardware inventory.	show chassis hardware on page 153
Routing Table Information	Information about the entries in the routing tables.	show route on page 516
Forwarding Table Information	Information about the entries in the kernel's forwarding table.	show route forwarding-table on page 543
IS-IS	Adjacent routers.	show isis adjacency on page 577
OSPF	Adjacent routers.	show ospf neighbor on page 607
BGP	Entries in the BGP neighbor database.	clear bgp neighbor on page 614
MPLS	Status of interfaces on which MPLS is running.	show mpls interface on page 698
	Configured LSPs on this router, as well as all ingress, transit, and egress LSPs.	show mpls lsp on page 699
	Routes that form a label-switched path.	show route label-switched-path on page 552
RSVP	Status of interfaces on which RSVP is running.	show rsvp interface on page 719
	Currently active RSVP sessions.	show rsvp session on page 727
	RSVP packet and error counters.	show rsvp statistics on page 734
Information for Customer Support	System information to collect before contacting customer support.	request support information on page 52

Craft Interface

You can use the craft interface to view and obtain status and troubleshooting information about the router. The craft interface is located on the front of the router and contains the following elements:

- System LEDs—Report the status of the Routing Engine, the status of each Flexible PIC Concentrator (FPC), and general system alarm conditions.
- System buttons—Used to reset clocks and stop alarms. Each FPC has a button that you press to take the FPC offline, allowing safe removal of the FPC.
- LCD display (on some routers)—Displays current system status or alarm conditions.
- Alarm relay contacts—Allow you to connect external alarm devices.
- Routing Engine ports—Allow you to connect the following external management devices:
  - Console port—Connects a system console to the Routing Engine with an RS-232 serial cable.
  - Auxiliary port—Connects a laptop or modem to the Routing Engine with an RS-232 serial cable.
  - Ethernet management port—Connects the Routing Engine to a management LAN (or any other device that plugs into an Ethernet connection) for out-of-band management of the router system. The Ethernet port can be 10 or 100 Mbps and uses an autosensing RJ-45 connector.

Part 2
System Management

- Monitor and Perform System Management Functions on page 41
- SNMP Monitoring and Troubleshooting on page 117
- M-Series Chassis Monitoring and Troubleshooting on page 127
- T-Series Platform Monitoring and Troubleshooting on page 167
- Accounting Options Monitoring and Troubleshooting on page 177

Chapter 3

Monitor and Perform System Management Functions

Table 6 summarizes the command-line interface (CLI) commands you can use to perform system management functions and monitor system management functions. In the table, the commands are grouped by functionality. In the remainder of this chapter, the commands are explained alphabetically.

Table 6: Commands for Monitoring and Performing System Management

Task	Task or Information to Monitor	Command
Software Version	Versions of software running on the router.	show version on page 109
	JUNOS software extensions.	show system software on page 97
Router Up Time	Current time and how long the router has been operational.	show system uptime on page 101
Configuration	Current running system configuration.	show configuration on page 67
	Verify the syntax of a configuration file.	test configuration on page 113
User Account Information	Users currently logged into the router.	show system users on page 102
	Login permissions for the current user.	show cli authorization on page 65
Log Files and Recent History	Log files and their contents and recent user logins.	show log on page 69
	Monitor the contents of log files.	monitor on page 47
	Monitor network traffic	ping on page 48
	Recent CLI history.	show cli history on page 66
Remote System Reachability	Check host reachability and network connectivity.	ping on page 48
	Check the reachability of a remote ATM node.	ping atm on page 50
	Check the reachability of a remote VPN node.	ping vpn-interface on page 51
	Determine the route to a network system.	traceroute on page 114
	Network connection information.	show system connections on page 89
Connecting to Remote Systems	Open secure shell connections between the local router and a remote system.	ssh on page 110
	Open a telnet session to a remote system.	telnet on page 112
Network Time Protocol	NTP peers and their state.	show ntp associations on page 70
	NTP peer variables.	show ntp status on page 71

Task	Task or Information to Monitor	Command
Packet Forwarding Engine	Packet Forwarding Engine FPC.	show pfe fpc on page 72
	Packet Forwarding Engine information.	show pfe terse on page 84
	Packet Forwarding Engine next-hop.	show pfe next-hop on page 73
	Packet Forwarding Engine route.	show pfe route on page 74
	Packet Forwarding Engine SCB.	show pfe scb on page 75
	Packet Forwarding Engine SFM.	show pfe sfm on page 77
	Packet Forwarding Engine statistics.	show pfe statistics dma on page 79 show pfe statistics error on page 80 show pfe statistics ip on page 81 show pfe statistics notification on page 82 show pfe statistics pio on page 82 show pfe statistics traffic on page 83
System Software	System memory and buffer usage information.	show system buffers on page 88
	Software processes running on the router.	show system processes on page 91
	Systemwide protocol-related statistics.	show system statistics on page 98
	Statistics about amount of free disk space in the router's file systems.	show system storage on page 100
	Routing, routing protocol, and interface tasks that are currently running.	show task on page 107
	JUNOS kernel memory usage.	show system virtual-memory on page 103
	State and checksum value for files in a file system.	show system audit on page 85
Hostname Lookup	Hostname lookup using DNS.	show host on page 68
ARP Table	Contents of the ARP table.	show arp on page 65
	Zero the contents of the ARP table.	clear arp on page 44
Stop and Restart Processes and Software	Restart a JUNOS software process.	restart on page 64
	Install software bundles or packages onto the router.	request system software add on page 59
	Remove software bundles or packages from the router.	request system software delete on page 61
	Rollback to previously installed version.	request system software rollback on page 62
	Check candidate software compatibility against the current configuration.	request system software validate on page 63
	Stop the routing software.	request system halt on page 53
	Reboot the routing software.	request system reboot on page 55
	Pending system halts or reboots.	show system reboot on page 97
	Clear a pending system halt or reboot.	clear system reboot on page 44
	Back up the file systems on the router.	request system snapshot on page 57
Manipulate Files	Copy a file to another location on the router or to another system on the network.	file copy on page 45
	List files and directories on the router.	file list on page 46
	Display the contents of a file.	file show on page 46
	Rename a file on the router.	file rename on page 46
	Delete a file on the router.	file delete on page 45

Task	Task or Information to Monitor	Command
Create a Shell	Create a UNIX-level shell.	start on page 111
Systemwide Messages	Send messages to users currently logged into the router.	request message on page 52
System Boot Messages	Display boot messages.	show system boot-messages on page 86
Information for Customer Support	Collect system information before contacting customer support.	request support information on page 52
Exit from CLI	Exit from the CLI to a UNIX shell.	quit on page 51

How to Specify Filenames and URLs

In some CLI commands and configuration statements—including file copy, load, save, set system login user *user-name* authentication *load-key-file*, and request system software add—you can include a filename. You can specify a filename or URL in one of the following ways:

- *filename*—File in the user's home directory (the current directory) on the local flash disk.
- *path*/*filename*—File on the local flash disk.
- /var/*filename* or /var/*path*/*filename*—File on the local hard disk.
- a:*filename* or a:*path*/*filename*—File on the local drive. The default path is / (the root-level directory). The removable media can be in MS-DOS or UNIX (UFS) format.
- *hostname:*/*path*/*filename, hostname:filename, hostname:path*/*filename,* or scp://*hostname*/*path*/*filename*—File on an scp/ssh client. This form is not available in the worldwide version of the JUNOS software. The default path is the user's home directory on the remote system. You can also specify *hostname* as *username@hostname*.
- ftp://*hostname*/*path*/*filename*—File on an FTP server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. The default path is the user's home directory. To specify an absolute path, the path must start with %2F; for example, ftp://*hostname*/%2F*path*/*filename*. To have the system prompt you for the password, specify prompt in place of the password. If a password is required, and you do not specify the password or prompt, an error message is displayed:

  user@host > **`file copy ftp://username@ftp.hostname.net//filename`**
  file copy ftp.*hostname*.net: Not logged in.
  user@host > **`file copy ftp://username:prompt@ftp.hostname.net//filename`**
  Password for *username*@ftp.*hostname*.net:

- http://*hostname*/*path*/*filename*—File on an HTTP server. You can also specify *hostname* as *username@hostname* or *username*:*password@hostname*. If a password is required and you omit it, you are prompted for it.
- re0:/*path*/*filename* or re1:/*path*/*filename*—File on a local Routing Engine.

clear arp

Syntax clear arp <hostname *hostname* > <vpn *vpn*>

Description Remove entries from the ARP table.

Options none—Clear all entries.

hostname *hostname*—(Optional) Name of specific host to clear.

vpn *vpn*—(Optional) Clear entries from the ARP table for specified VPN.

Required Privilege Level clear

Sample Output

```
user@host> clear arp
user@host>
```

clear log

Syntax clear log *(filename* | all)

Description Remove contents of a log file.

Options all—Truncate the current log file and delete all archived log files.

filename—Name of the specific log file to truncate.

Required Privilege Level clear

clear system reboot

Syntax clear system reboot

Description Clear any pending system software reboots or halts.

Required Privilege Level maintenance

See Also request system reboot on page 55

Sample Output

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

configure exclusive

Syntax configure exclusive

Description Exclude others from modifying the configuration database. Attempts by other users to modify the configuration while configure exclusive is active result in a descriptive error message, giving the owner, tty, and timestamp of the exclusive lock.

Required Privilege Level configure

See Also request system logout user on page 54

file copy

Syntax file copy *source destination*

Description Copy files from one place to another on the local router or between the local router and a remote system.

Options *destination*—Destination of the copied file. Specify this as a URL or filename, as described in "How to Specify Filenames and URLs" on page 43. If you are copying a file to the current directory (your home directory on the local router) and are not renaming the file, specify the destination with a period (.).

source—Source of the original file. Specify this as a URL or filename, as described in "How to Specify Filenames and URLs" on page 43.

Required Privilege Level maintenance

Sample Output

```
user@host> copy /var/tmp/rpd.core.4 berry:/c/junipero/tmp

...transferring.file...... |          0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
```

The following is an example of copying a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

file delete

Syntax file delete *filename*

Description Delete a file on the local router.

Options *filename*—Name of the file to delete.

Required Privilege Level maintenance

Sample Output

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax file list *filename*

Description Display a list of files on the local router.

Options *filename*—Name of a specific directory or file.

Required Privilege Level maintenance

Sample Output

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file rename

Syntax file rename *source destination*

Description Rename a file on the local router.

Options *destination*—New name for the file.

source—Original name of the file.

Required Privilege Level maintenance

Sample Output

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file show

Syntax file show *filename*

Description Display the contents of a file.

Options *filename*—Name of a file.

Required Privilege Level maintenance

Sample Output

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
Apr 13 21:07:36 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:08:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
...
```

monitor

Syntax monitor (start | stop) *list*

Description Display the end of system log or trace files and additional entries being added to those files.

Log files are generated by the routing protocol process or by syslog. The log files generated by syslog are those configured with the syslog statement at the [edit system] hierarchy level and the options statement at the [edit routing-options] hierarchy level. The trace files generated by the routing protocol process are those configured with traceoptions statements at the [edit routing-options], [edit interfaces], and [edit protocols *protocol*] hierarchy levels.

For explanation of specific errors, see the show interfaces extensive command.

Options *list*—Display status of monitored files.

start—Start displaying the file contents.

stop—Stop displaying the file contents.

Required Privilege Level trace

See Also show interfaces on page 194 and monitor traffic on page 452

Output Fields **** filename ****—Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.

Sample Output

```
user@host> monitor start system-log
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from trip.jcmax.com
Jul 20 15:07:37 hang sshd[5847]: log: executing remote command as root: scp -t /
tmp
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

ping

Syntax ping *host* < interface *source-interface*> <bypass-routing> <count *requests*> <do-not-fragment> <interval *seconds*> <pattern *string*> <record-route> <routing-instance *routing-instance-name*> <size *bytes*> <strict> <tos *type-of-service*> <ttl *value*> <via *route*> <rapid | detail>

Description Check host reachability and network connectivity.

The ping command sends ICMP ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host.

Type Ctrl-C to interrupt a ping command.

See Also ping atm on page 50 and ping vpn-interface on page 51

Options bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

count *requests*—(Optional) Number of ping requests to send.
Range: 0 through 1,000,000
Default: No limit to the number of requests sent. Type Ctrl-C to interrupt a ping command.

detail—(Optional) Include in the output the interface on which the ping reply was received.

do-not-fragment—(Optional) Set the do-not-fragment (DF) bit in the IP header of the ping packets.

host—IP address or name of the remote system to ping.

interface *source-interface*—(Optional) Interface to use to send the ping requests.

interval *seconds*—(Optional) How often to send ping requests.
Range: 0.1 through 10,000 seconds
Default: 1 second

pattern *string*—(Optional) Specify a 16-bit pattern string to include in the ping packet.
Range: 0 through 65536

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the ping attempt.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of request, include the count option.

record-route—(Optional) Record and report the packet's path.

size *bytes*—(Optional) Size of ping request packets.
Range: 0 through 65468 bytes
Default: 56 bytes, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

strict—(Optional) Use strict source route option.

tos *type-of-service*—(Optional) Set the type-of-service (tos) bit in the IP header of the ping packets.
Range: 0 through 255

ttl *value*—(Optional) TTL value to include in the ping request.
Range: 0 through 255

via *route*—(Optional) Specify the intermediate router through which the packet should traverse on its way to the system being pinged. Configuring an intermediate router sets the loose source route field in the IP header.

Required Privilege Level network

Sample Output

```
user@host> ping santacruz
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]

user@host> ping santacruz size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms

user@host> ping santacruz rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping atm

Syntax ping atm interface *interface* <count *count*> <end-to-end | segment> <interval *interval*> <sequence-number *sequence-number*> <vci *vci*> <brief>

Description Check the reachability of a remote ATM node. All packets are 53 bytes.

Type Ctrl-C to interrupt a ping atm command.

Options brief—(Optional) Display only the ATM ping summary statistics. These are displayed after you type Ctrl-C to interrupt the ping atm command.

count *count*—(Optional) Number of ping requests to send.
Range: 0 through 10,000
Default: No limit to the number of requests sent. Type Ctrl-C to interrupt a ping command.

end-to-end—(Optional) Cells are sent to the end node. This is the default.

interface *interface*—Interface to use to send the ATM ping requests.

interval *interval*—(Optional) How often to send ping requests.
Range: 0.1 through 10,000 seconds
Default: 1 second

segment—(Optional) Cells are sent only to the intermediate node.

sequence-number *sequence-number*—(Optional) Starting sequence number (correlation tag).
Range: 0 through 65468 bytes
Default: 1

vci *vci*—(Optional) ATM point-to-point virtual circuit identifier. It can be a virtual circuit identifier (vci) or a virtual private identifier (vpi.vci).

Required Privilege Level network

Sample Output

```
user@host> ping atm vci 0.128 interface at-5/3/0
53 byte oam cell received on (vpi=0 vci=128): seq=1
53 byte oam cell received on (vpi=0 vci=128): seq=2
53 byte oam cell received on (vpi=0 vci=128): seq=3
53 byte oam cell received on (vpi=0 vci=128): seq=4
53 byte oam cell received on (vpi=0 vci=128): seq=5
^C[abort]
--- atmping statistics ---
5 cells transmitted, 5 cells received, 0% cell loss

user@host> ping atm vci 0.128 interface at-5/3/0 brief
^C[abort]
--- atmping statistics ---
9 cells transmitted, 9 cells received, 0% cell loss
```

ping vpn-interface

Syntax ping vpn-interface *vpn-interface host* <local *echo-address*>

Description Check the reachability of a remote vpn-interface node.

The ping vpn-interface command provides the source address of the ICMP packets and refers to the appropriate VPN routing and forwarding (VRF) table to perform a route lookup.

Type Ctrl-C to interrupt a ping vpn-interface command.

Options *host*—IP address or name of remote system to ping.

local *echo-address*—(Optional) Set the source IP address (the local address) in outgoing ECHO_REQUEST packets.
If default-address-selection is configured, lo0 is the source.
If default-address-selection is not configured, outgoing interface is the source.

vpn-interface—Name of VPN interface from which to ping. Specify address including the logical interface.
If no logical interface is specified, lo0 is used.

Sample Output

```
user@vpnhost> ping vpn-interface fe-1/1/0.1 192.168.197.94 local
192.168.197.93 count 3
PING 192.168.197.94 (192.168.197.94): 56 data bytes
64 bytes from 192.168.197.94: icmp_seq=0 ttl=255 time=0.866 ms
64 bytes from 192.168.197.94: icmp_seq=1 ttl=255 time=0.728 ms
64 bytes from 192.168.197.94: icmp_seq=2 ttl=255 time=0.753 ms
--- 192.168.197.94 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.728/0.782/0.866/0.060 ms
```

quit

Syntax quit

Description Exit from the command-line interface (CLI). Issuing this command places you into a UNIX shell. To restart the CLI, type cli.

Required Privilege Level Users with any privilege can issue this command.

request message

Syntax request message all message "*text*"
request message message "*text*" (terminal *terminal-name* | user *user-name*)

Description Display a message on the screens of all users who are logged into the router or on specific screens.

Options all—Display a message on the terminal of all users who are currently logged in.

message "*text*"—Message to display.

terminal *terminal-name*—Name of the terminal on which to display the message.

user *user-name*—Name of the user to whom to direct the message.

Required Privilege Level maintenance

Sample Output

```
user@host> request message message “Maintenance window in 10 minutes” user maria
Message from user@host on ttyp0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

request support information

Syntax request support information

Description Display information about the system. You should issue this command before contacting customer support and then include the command output in your support request.

The output of this command is generally quite long, so you might want to redirect the output to a file.

This command is a combination of the following CLI operational mode commands:

- show version
- show chassis firmware
- show chassis hardware
- show chassis environment
- show interfaces extensive for each configured interface
- show configuration (excluding any SECRET-DATA)

Required Privilege Level maintenance

Sample Output

```
user@host> request support information | save filename
Wrote 1143 lines of output to ‘filename’
user@host>
```

request system halt

Syntax request system halt <at *time*> <in *minutes*> <message "*text*">

Description Stop the router software.

Options at *time*—(Optional) Time at which to stop the software. You can specify time in one of the following ways:

- now—Stop the software immediately. This is the default.
- *+minutes*—Number of minutes from now to stop the software.
- *yymmddhhmm*—Absolute time at which to stop the software, specified as the year, month, day, hour, and minute.
- *hh:mm*—Absolute time on the current day at which to stop the software, specified in 24-hour time, using a colon to separate hours from minutes.

in *minutes*—(Optional) Number of minutes from now to stop the software. This option is an alias for the at *+minutes* option.

message "*text*"—(Optional) Message to display to all system users before stopping the software.

Default: at now (stop the software immediately)

Required Privilege Level maintenance

Examples Assume that it is currently 5 p.m.

Stop the system in two hours (at 7 p.m.):

```
request system halt at +120
request system halt in 120
request system halt at 19:00
```

Stop the system immediately:

```
request system halt at now
```

Stop the system at 1:20 a.m. Because 1:20 a.m. is the next day, you must specify the absolute time.

```
request system halt at yymmdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

Sample Output

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes
*** FINAL System shutdown message from root@lab8 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done

The operating system has halted.
Please press any key to reboot.
```

request system logout user

Syntax request system logout user *user* <terminal *terminal* | pid *pid*>

Description Log out a specific user from the router and the configuration database. If the user held the configure exclusive lock, clears the exclusive lock.

Options *user*—Name of user to logout.

pid *pid*—(Optional) Management (MGD) process identifier for user.

terminal *terminal*—(Optional) Name of terminal to log out. Logs out the named user on the named terminal session. If no terminal is named, logs out the user on all terminal sessions.

Required Privilege configure

See Also configure exclusive on page 44

request system reboot

Syntax request system reboot <at *time*> <in *minutes*> <message "*text*">

Description Reboot the software.

Reboot requests are recorded to the system log files, which you can view with the show log command. Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the show system processes command.

Options at *time*—(Optional) Time at which to reboot the software. You can specify time in one of the following ways:

- now—Stop or reboot the software immediately. This is the default.
- *+minutes*—Number of minutes from now to reboot the software.
- *yymmddhhmm*—Absolute time at which to reboot the software, specified as the year, month, day, hour, and minute.
- *hh:mm*—Absolute time on the current day at which to stop the software, specified in 24-hour time, using a colon to separate hours from minutes.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the at *+minutes* option.

message "*text*"—(Optional) Message to display to all system users before stopping or rebooting the software.

Default: at now (reboot the system immediately)

Required Privilege Level maintenance

See Also clear system reboot on page 44

Examples Assume that it is currently 5 p.m.

Reboot the system in two hours (at 7 p.m.):

```
request system reboot at +120
request system reboot in 120
request system reboot at 19:00
```

Reboot the system immediately:

```
request system reboot at now
```

Reboot the system at 1:20 a.m. Because 1:20 a.m. is the next day, you must specify the absolute time.

```
request system reboot at yymmdd120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

Sample Output

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
user@host>
user@host> request system reboot at 2300 message “Maintenance time!”
Reboot the system ? [yes,no] (no) yes
shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system snapshot

Syntax request system snapshot <partition>

Description Back up the currently running and active file system partitions on the router to standby partitions that are not running. Specifically, the root file system (/) is backed up to /altroot and /config is backed up to /altconfig. The root and /config file systems are on the router's flash drive, and the /altroot and /altconfig file systems are on the router's hard drive.

> **Note**
>
> Prior to upgrading the software on the router, when you have a known stable system, you should issue the request system snapshot command to back up the software, including the configuration, onto the /altroot and /altconfig file systems.
>
> After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, you should issue the request system snapshot command to back up the new software onto the /altroot and /altconfig file systems.
>
> After you run the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

You cannot issue the request system snapshot command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Options partition—(Optional) Repartition the flash drive before a snapshot occurs. If the partition table on the flash drive is corrupted, the request system snapshot command fails and reports errors.

The partition option is only supported for restoring the software image from the hard drive to the flash drive. You cannot issue the request system snapshot command when you enable flash disk mirroring. We recommend that you disable flash disk mirroring when you upgrade or downgrade the software. For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Required Privilege Level maintenance

Output Fields unmount—Unmounting a file system.

Sample Output: request system snapshot

```
user@host> request system snapshot
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)
umount: /altconfig: not currently mounted
Copying /config to /altconfig.. (this may take a few minutes)

The following filesystems were archived: / /config
```

Sample Output: request system snapshot when partition flag is on

```
user@host>request system snapshot partition
Performing preliminary partition checks ...
Partitioning ad0 ...
umount: /altroot: not currently mounted
Copying / to /altroot.. (this may take a few minutes)

The following filesystems were archived: / /config
```

Sample Output: request system snapshot when mirroring is enabled

```
user@host> request system snapshot
Snapshot is not possible since mirror-flash-on-disk is configured.
```

request system software add

Syntax request system software add *package-name* < no-copy > < delay-restart> <reboot > < validate >

Description Install a software bundle or package onto the router.

> **Note**
>
> Prior to upgrading the software on the router, when you have a known stable system, you should issue the request system snapshot command to back up the software, including the configuration, onto the /altroot and /altconfig file systems.
>
> After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, you should issue the request system snapshot command to back up the new software onto the /altroot and /altconfig file systems.
>
> After you run the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, jkernel, last. Add the operating system package, jkernel, first and the routing software package, jroute, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernel
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

Options delay-restart—(Optional) Install software bundle or package but do not restart daemons.

no-copy—(Optional) Install software bundle or package but do not save copies of package files.

package-name—Name of the software bundle or package to install or remove. Specify the filename as described in "How to Specify Filenames and URLs" on page 43.

reboot—(Optional) After adding the software bundle or package, reboot the system.

validate—(Optional) Validate the package name against the current configuration as a prerequisite to adding the software.

Required Privilege Level maintenance

See Also request system software delete on page 61, request system software rollback on page 62, and request system software validate on page 63

Sample Output: request system software add (validate), successful case

```
user@host> request system software add validate
/c/sjg/work/current/ship/route-5.3I20020124_2231_sjg-signed.tgz

Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /c/sjg/work/current/ship/jroute-5.3I20020124_2231_sjg-signed.tgz
Verified MD5 checksum of
/var/chroot/var/tmp/jroute/jroute-5.3I20020124_2231_sjg.tgz
Using /var/chroot/var/tmp/jroute-signed/jroute-5.3I20020124_2231_sjg.tgz
Using /packages/jkernel-5.3I20020124_0520_sjg
Validating against /config/juniper.conf.gz
mgd: commit complete
Installing package
'/c/sjg/work/current/ship/jroute-5.3I20020124_2231_sjg-signed.tgz' ...
Verified MD5 checksum of jroute-5.3I20020124_2231_sjg.tgz
Auto-deleting old jroute...
Adding jroute...
Mounted jroute package on /dev/vn6...
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting rpd ...
Restarting ppmd ...
Restarting aprobed ...
Saving package file in /var/sw/pkg/jroute-5.3I20020124_2231_sjg-signed.tgz ...
Saving state for rollback ...

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

Sample Output: request system software add (validate), failure case

```
user@host> request system software add validate
/var/sw/pkg/jroute-5.0I20020111_2246_sjg.tgz

Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jroute-5.0I20020111_2246_sjg.tgz
WARNING: Package 'jroute' is not compatible with package 'jkernel':
WARNING:     Routing socket version mismatch
WARNING:     (jroute:16 != jkernel:17)
WARNING: Package 'jroute' is not compatible with package 'jkernel':
WARNING:     The configuration database layout is incompatible
WARNING:     (jroute:2 != jkernel:3)
ERROR: validate-config: jroute/+REQUIRE fails
WARNING: Current configuration not compatible with
/var/sw/pkg/jroute-5.0I20020111_2246_sjg.tgz
```

request system software delete

Syntax request system software delete *package-name* <jcrypto | jdocs | jkernel | jpfe | jroute | junos>

Description Remove a software bundle or package from the router.

> **Note**
>
> Prior to upgrading the software on the router, when you have a known stable system, you should issue the request system snapshot command to back up the software, including the configuration, onto the /altroot and /altconfig file systems.
>
> After you have upgraded the software on the router and are satisfied that the new packages are successfully installed and running, you should issue the request system snapshot command to back up the new software onto the /altroot and /altconfig file systems.
>
> After you run the request system snapshot command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Options jcrypto—(Optional, in domestic version only) JUNOS security software.

jdocs—(Optional) JUNOS online documentation file.

jkernel—(Optional) JUNOS kernel software suite.

jpfe—(Optional) JUNOS packet forwarding engine support.

jroute—(Optional) JUNOS routing software suite.

junos—(Optional) JUNOS base software.

package-name—Name of the software bundle or package to install or remove. Specify the filename as described in "How to Specify Filenames and URLs" on page 43.

Required Privilege Level maintenance

See Also request system software add on page 59, request system software rollback on page 62, and request system software validate on page 63

request system software rollback

Syntax request system software rollback

Description Revert to the set of software as of the last successful request system software add. Rollback fails if any required package (or a jbundle package containing the required package) cannot be found in /var/sw/pkg.

> **Note**
>
> You cannot issue the request system software rollback command to return to the previously installed software after using a jinstall package.
>
> To return to the previously installed software, use the jinstall package that corresponds with the previously installed software.

Required Privilege Level maintenance

See Also request system software add on page 59, request system software delete on page 61, and request system software validate on page 63

request system software validate

Syntax request system software validate *package-name*

Description Validate candidate software against the current configuration of the router.

Test if the candidate software is compatible with the current configuration. The candidate software is loaded and committed. If all daemons validate the configuration successfully, then the candidate software is compatible. If not, an error message is generated.

Required Privilege Level maintenance

See Also request system software add on page 59, request system software delete on page 61, and request system software rollback on page 62

Sample Output: request system software validate, successful case

```
user@host> request system software validate
/var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

Sample Output: request system software validate, failure case

```
user@host>request system software validate
/var/sw/pkg/jbundle-5.0I20020110_0958_sjg.tgz

Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.0I20020110_0958_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.0I20020110_0958_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: rtslib: FATAL ERROR interface version mismatch: kernel=18 library=16,a
reboot or software upgrade may be required
mgd:
mgd: error: configuration check-out failed
WARNING: Current configuration not compatible
with/var/sw/pkg/jbundle-5.0I20020110_0958_sjg.tgz
```

restart

Syntax restart (class-of-service | interface-control | mib-process | network-access-service | remote-operations | routing | sampling | snmp) <gracefully> <immediately> <soft>

Description Restart a JUNOS software process.

Options class-of-service—(Optional) Restart the class of service process, which controls the router's class of service configuration.

gracefully—(Optional) Restart the software process by sending the equivalent of a UNIX SIGTERM signal.

immediately—(Optional) Immediately restart the process by sending the equivalent of a UNIX SIGKILL signal.

interface-control—(Optional) Restart the interface process, which controls the router's physical interface devices and logical interfaces.

mib-process—(Optional) Restart the MIB II process, which provides the router's MIB II agent.

network-access-service—(Optional) Restart the network access process, which provides the router's CHAP authentication service.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

routing—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router and maintains the routing tables.

sampling—(Optional) Restart the sampling process, which performs packet sampling and cflowd export.

snmp—(Optional) Restart the SNMP process, which provides the router's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. This option is the equivalent of a UNIX SIGHUP signal; omitting this option is the equivalent of a UNIX SIGTERM (kill) operation.

> **Note** Never restart any of the software processes unless instructed to do so by a customer support engineer.

Default: gracefully

Required Privilege Level reset

Sample Output

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
user@host>
```

show arp

Syntax show arp <no-resolve>

Description Display the entries in the ARP table.

Options no-resolve—(Optional) Do not attempt to determine the host name that corresponds to the IP address.

Required Privilege Level view

Sample Output

```
user@host> show arp
MAC Address       Address         Name                      Flags
00:00:0c:06:2c:0d 192.168.1.2     firewall.my.net
```

show cli authorization

Syntax show cli authorization

Description Display the permissions for the current user.

Required Privilege Level view

Sample Output

```
root@host> show cli authorization
Current user: 'root' login: ‘boojum’ class '(root)'
Permissions:
    admin               -- Can view user accounts
    admin-control       -- Can modify user accounts
    clear               -- Can clear learned network information
    configure           -- Can enter configuration mode
    control             -- Can modify any config
    edit                -- Can edit full files
    field               -- Special for field (debug) support
    floppy              -- Can read and write from the floppy
    interface           -- Can view interface config
    interface-control   -- Can modify interface config
    network             -- Can access the network
    reset               -- Can reset/restart interfaces and daemons
    routing             -- Can view routing config
    routing-control     -- Can modify routing config
    shell               -- Can start a local shell
    snmp                -- Can view SNMP config
    snmp-control        -- Can modify SNMP config
    system              -- Can view system config
    system-control      -- Can modify system config
    trace               -- Can view trace file settings
    trace-control       -- Can modify trace file settings
    view                -- Can view current values and statistics
    maintenance         -- Can become the super-user
    firewall            -- Can view firewall config
    firewall-control    -- Can modify firewall config
    secret              -- Can view secret config
    secret-control      -- Can modify secret config
```

show cli history

Syntax show cli history <*count*>

Description Display list of previous CLI commands

Options *count*—(Optional) Maximum number of commands to display.

Required Privilege Level view

Sample Output

```
user@host> show cli history
  11:14:14 -- show arp
  11:22:10 -- show cli authorization
  11:27:12 -- show cli history
```

show configuration

Syntax show configuration <*statement-path*>

Description Display the configuration that currently is running on the router, which is the last committed configuration. If you have modified the configuration since you last committed it, the configuration information displayed by the show configuration command will be different from that displayed by issuing the show command from the [edit] hierarchy level in configuration mode.

The portions of the configuration that you are allowed to see depend on the user class you are a member of and the permissions that the user class has. If you do not have permission to view a portion of the configuration, the text ACCESS-DENIED is substituted in that portion of the configuration. If you do not have permission to view authentication keys and passwords in the configuration (that is, if the secret permission bit is not set for your user account), the text SECRET-DATA is substituted in that portion of the configuration.

If an identifier in the configuration contains a space, the identifier is shown in quotation marks when it is displayed.

Options none—Display the entire configuration.

<*statement-path*>—(Optional) Display a particular hierarchy in a configuration.

Required Privilege Level view

Sample Output: show configuration

```
user@host> show configuration
system {
    host-name eagle;
    domain-name east.net;
    backup-router 192.1.1.254;
    time-zone America/Los_Angeles;
    debugger-on-panic;
    debugger-on-break;
    default-address-selection;
    dump-on-panic;
    name-server {
        192.154.169.254;
        192.154.169.249;
        192.154.169.176;
    }
    services {
        telnet;
    }
    tacplus-server {
        1.2.3.4 {
            secret /* SECRET-DATA */;
            ...
        }
    }
}
```

```
interfaces {
...
protocols {
    isis {
        export "direct routes";
    }
}
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}
```

Sample Output: show configuration policy-options

```
user@host> show configuration policy-options
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}
```

show host

Syntax show host *host-name*

Description Display DNS hostname information about a particular host.

Options *host-name*—Hostname or address.

Required Privilege Level view

Sample Output

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254
user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

show log

Syntax show log <user <*user-name*>> <*filename*>

Description List log files, display log file contents, and display information about users who have logged into the router.

Options none—List all log files.

filename—(Optional) Display the log messages in the specified log file.

user <*user-name*>—(Optional) Display logging information about users who have recently logged into the router. If you include *user-name*, display logging information about the specified user.

Required Privilege Level trace

Sample Output

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
-rw-rw-r--  1 root  bin          27 Aug 26 18:00 wtmp.0.gz
-rw-rw-r--  1 root  bin          46 Aug 20 01:04 wtmp.1.gz
-rw-rw-r--  1 root  bin          63 Aug 19 18:28 wtmp.2.gz

user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to “/var/log/rpd” started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint
Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...

user@host> show log user
darius   mg2546                    Thu Oct  1 19:37   still logged in
darius   mg2529                    Thu Oct  1 19:08 - 19:36  (00:28)
darius   mg2518                    Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                    Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    jun.site.per    Wed Sep 30 18:39 - 18:41  (00:02)
alex     ttyp1    192.168.1.2      Wed Sep 30 01:03 - 01:22  (00:19)
...
```

show ntp associations

Syntax show ntp associations

Description Display NTP peers and their state. This command is equivalent to the UNIX ntpq –c peers command.

Required Privilege Level view

Output Fields remote—Address or name of the remote NTP peer.

refid—Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of 0.0.0.0.

st—Stratum of the remote peer.

t—Type of peer. It can be b (broadcast), l (local), m (multicast), or u (unicast).

when—When the last packet from the peer was received.

poll—Polling interval, in seconds.

reach—Reachability register, in octal.

delay—Current estimated delay of the peer, in milliseconds.

offset—Current estimated offset of the peer, in milliseconds.

disp—Current estimated dispersion of the peer, in milliseconds.

Character in the left margin—Indicates the fate of the peer in the clock selection process. It can be one of the following:

- space—Discarded because of a high stratum value or failed sanity checks.
- x—Designated falseticker by the intersection algorithm.
- . —Culled from the end of the candidate list.
- – —Discarded by the clustering algorithm.
- +—Included in the final selection set.
- #—Selected for synchronization, but distance exceeds maximum.
- *—Selected for synchronization.
- o—Selected for synchronization, but pps signal is in use.

Sample Output

```
user@host> show ntp associations
     remote           refid      st t when poll reach   delay   offset    disp
==============================================================================
*wolfe-gw.junipe tick.ucla.edu    2 u   43   64  377     1.86    0.319    0.08
```

show ntp status

Syntax show ntp status

Description Display the values of internal variables returned by NTP peers. This command is equivalent to the UNIX ntpq –c readlist command.

Required Privilege Level view

Sample Output

```
user@host> show ntp status
status=06f4 leap_none, sync_ntp, 15 events, event_peer/strat_chg
system="FreeBSD", leap=00, stratum=4, rootdelay=7.90,
rootdispersion=22.99, peer=23732, refid=olive-oyl.east.net,
reftime=b97c36cf.bed61000  Wed, Aug 12 1998 15:44:15.745, poll=7,
clock=b97c3715.ba307000  Wed, Aug 12 1998 15:45:25.727, phase=-0.283,
freq=-23633.96, error=0.695
```

show pfe fpc

Syntax show pfe fpc <*slot*>

Description Display statistics from a particular FPC slot.

Options *slot*—(Optional) Slot number. It can be a value from 0 through 7.

Required Privilege Level admin

Sample Output

```
user@host> show pfe fpc 1
FPC 1 status:
  Slot:                 Present
  State:                Online
  Last State Change:    2000-01-10 18:12:27 UTC
  Uptime:               1d 03:31
  Failures:             0
  Pending:              0

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                0
  Wakeup:               0
  Resync Request:       0
  Resync Done:          0
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface        2251         2219
     Chassis           0            0
        Boot           0            0
    Next-hop           0            0
       Jtree           0            0
       Cprod           0            0
       Route           0            0
         Pfe           0            1
         Dfw           0            0
  Mastership           0            0
       Empty           0            0

PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0           52273
     1               0
     2               0
    …
    21               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0           52273
     1               0
…
    21               0
```

show pfe next-hop

Syntax show pfe next-hop <interface> <sfm *slot*>

Description Display information about the Packet Forwarding Engine next hop.

Options interface—(Optional) Display Packet Forwarding Engine next-hop interface.

sfm *slot*—(Optional, for systems with multiple Switching and Forwarding Modules only.) Show the next hops for an SFM slot. *slot* can be a value from 0 through 3.

Required Privilege Level admin

Sample Output

```
user@host> show pfe next-hop

SFM 0

Nexthop Info:
   ID      Type      Interface   Protocol         Encap  Next Hop Addr    MTU
-----  --------  -------------  ---------  ------------  ---------------  -----
    4     Mcast  -                   IPv4             -  0.0.0.0              0
    5     Bcast  -                   IPv4             -  -                    0
    7   Discard  -                   IPv4             -  -                    0
    8  MDiscard  -                   IPv4             -  -                    0
    9    Reject  -                   IPv4             -  -                    0
   13     Local  -                   IPv4             -  192.168.4.60         0
   14   Resolve  fxp0.0              IPv4   Unspecified  -                    0
   17     Local  -                   IPv4             -  127.0.0.1            0
   18   Unicast  fxp0.0              IPv4   Unspecified  192.168.4.254        0
   21     Local  -                   IPv4             -  11.1.0.1             0
   22   Unicast  at-0/1/0.0          IPv4      ATM SNAP  11.1.0.2          4482
   …
   26   Unicast  at-0/3/0.0          IPv4      ATM SNAP  10.3.0.2          4482

SFM 1

Nexthop Info:
   ID      Type      Interface   Protocol         Encap  Next Hop Addr    MTU
-----  --------  -------------  ---------  ------------  ---------------  -----
    4     Mcast  -                   IPv4             -  0.0.0.0              0
    5     Bcast  -                   IPv4             -  -                    0
    7   Discard  -                   IPv4             -  -                    0
    8  MDiscard  -                   IPv4             -  -                    0
    9    Reject  -                   IPv4             -  -                    0
   13     Local  -                   IPv4             -  192.168.4.60         0
   14   Resolve  fxp0.0              IPv4   Unspecified  -                    0
   17     Local  -                   IPv4             -  127.0.0.1            0
   18   Unicast  fxp0.0              IPv4   Unspecified  192.168.4.254        0
   21     Local  -                   IPv4             -  11.1.0.1             0
   22   Unicast  at-0/1/0.0          IPv4      ATM SNAP  11.1.0.2          4482
   23     Local  -                   IPv4             -  10.2.0.1             0
   24   Unicast  at-0/2/0.0          IPv4      ATM SNAP  10.2.0.2          4482
   25     Local  -                   IPv4             -  10.3.0.1             0
   26   Unicast  at-0/3/0.0          IPv4      ATM SNAP  10.3.0.2          4482
```

show pfe route

Syntax show pfe route <ip | mpls | statistics> <sfm *slot*> <summary>

Description Display the Packet Forwarding Engine routing table.

Options ip—(Optional) Display Packet Forwarding Engine IPv4 routes.

mpls—(Optional) Display Packet Forwarding Engine MPLS information.

sfm *slot*—(Optional, for systems with multiple Switching and Forwarding Modules only.) Display routes for an SFM slot. *slot* can be a value from 0 through 3.

statistics—(Optional) Display Packet Forwarding Engine statistics.

summary—(Optional) Display summary of Packet Forwarding Engine information.

Required Privilege Level admin

show pfe scb

Syntax show pfe scb

Description (M40 routers only) Display information about the Packet Forwarding Engine SCB.

Required Privilege Level admin

Sample Output

```
user@host> show pfe scb
SCB status:
  Slot:                 Present
  State:                Online
  Last State Change:    1999-02-05 11:02:36 UTC
  Uptime:               1d 02:31
  Failures:             0
  Pending:              0

PFE listener statistics:
  Open:                 1
  Close:                0
  Sleep:                1
  Wakeup:               0
  Resync Request:       1
  Resync Done:          1
  Resync Fail:          0
  Resync Time:          0

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface       10715        10594
     Chassis           0            0
        Boot           0            0
    Next-hop           8            0
       Jtree           0            0
       Cprod           0            0
       Route          11            0
         Pfe        1592         1593
         Dfw           0            0
  Mastership           0            0
       Empty           0            0
```

```
PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0            5298
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               3
    12               3
    13               1
    14               3
    15               7
    16              41
    17               0
    18               0
    19               0
    20               0
    21               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0            5298
     1               0
     2               0
     3               0
     4               2
     5               3
     6               1
     7               1
     8               1
     9               6
    10              44
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0
```

show pfe sfm

Syntax show pfe sfm *slot*

Description For systems with multiple Switching and Forwarding Modules only, display information about an SFM in the Packet Forwarding Engine SFM.

Options *slot*—Display statistics from the specified SFM slot. *slot* can be 0, 1, 2, or 3.

Required Privilege Level admin

Sample Output

```
user@host> show pfe sfm 1
SFM 1 status:
  Slot:                 Offline
  State:                Init
  Last State Change:    2000-03-01 07:45:55 UTC
  Downtime:             17:47:29
  Failures:             167
  Pending:              0

PFE listener statistics:
  Open:                 167
  Close:                167
  Sleep:                2
  Wakeup:               1
  Resync Request:       2
  Resync Done:          2
  Resync Fail:          0
  Resync Time:          1

PFE IPC statistics:
 type        TX Messages  RX messages
 ----------- -----------  -----------
      Header           0            0
        Test           0            0
   Interface           0            0
     Chassis           0            0
        Boot           0            0
    Next-hop           0            0
       Jtree           0            0
       Cprod           0            0
       Route           0            0
         Pfe           0            0
         Dfw           0            0
  Mastership           0            0
       Empty           0            0
```

```
PFE socket-buffer mbuf depth:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0

PFE socket-buffer bytes pending transmit:
 bucket          count
 ------          -----
     0               0
     1               0
     2               0
     3               0
     4               0
     5               0
     6               0
     7               0
     8               0
     9               0
    10               0
    11               0
    12               0
    13               0
    14               0
    15               0
    16               0
    17               0
    18               0
    19               0
    20               0
    21               0
```

show pfe statistics dma

Syntax show pfe statistics dma <sfm *slot*>

Description Display statistics about the Packet Forwarding Engine direct memory access.

Options sfm *slot*—(Optional, for systems with multiple Switching and Forwarding Modules only.) Display routes for an SFM slot. *slot* can be a value from 0 through 3.

Required Privilege Level admin

Sample Output

```
user@host> show pfe statistics dma

DMA Statistics:

           Name    Requests   Completed      Failed
 --------------  ----------  ----------  ----------
    Packet Read      905119      905119           0
   Packet Write      943761      943761           0
  Physical Read           0           0           0
 Physical Write           0           0           0

DMA Errors:

           Name     Write 0     Write 1      Read 0      Read 1
 --------------  ----------  ----------  ----------  ----------
   Illegal Bank           0           0           0           0
  Address Range           0           0           0           0
      ECC Error           0           0           0           0
    PCI Retries           0           0           0           0
      PCI Error           0           0           0           0

DMA Requests:
 Requests available: 256, Requests used: 0
```

show pfe statistics error

Syntax show pfe statistics error

Description Display statistics about the Packet Forwarding Engine errors.

Required Privilege Level admin

Sample Output

```
user@host> show pfe statistics error

PFE error statistics:
        C chip    A1 chip    A2 chip
    ---------- ---------- ----------
             0          0          0  scan fail
             0          0        N/A  A1<->C FCS error
             0        N/A          0  A2<->C FCS error
           N/A          0          0  A<->B FCS error
B chip slots:
             0          1          2          3
    ---------- ---------- ---------- ----------
             0          0          0          0   scan fail
             0          0          0          0   A1->B FCS error
             0          0          0          0   A2->B FCS error
             0          0          0          0   correctable ECC error
             0          0          0          0   uncorrectable ECC error
             0          0          0          0   multiple ECC errors
             0          0          0          0   B->HS link error
             0          0          0          0   A1->Bm error
             0          0          0          0   A2->Bo error
             0          0          0          0   write buffer overflow
             0          0          0          0   Bo FIFO sync error
             0          0          0          0   Bo FIFO size error
             0          0          0          0   Bo stream stuck error
             0          0          0          0   Bo SRAM parity error

             4          5          6          7
    ---------- ---------- ---------- ----------
             0          0          0          0   scan fail
             0          0          0          0   A1->B FCS error
             0          0          0          0   A2->B FCS error
             0          0          0          0   correctable ECC error
             0          0          0          0   uncorrectable ECC error
             0          0          0          0   multiple ECC errors
             0          0          0          0   B->HS link error
             0          0          0          0   A1->Bm error
             0          0          0          0   A2->Bo error
             0          0          0          0   write buffer overflow
             0          0          0          0   Bo FIFO sync error
             0          0          0          0   Bo FIFO size error
             0          0          0          0   Bo stream stuck error
             0          0          0          0   Bo SRAM parity error
```

show pfe statistics ip

Syntax show pfe statistics ip <icmp | options> <sfm *slot*>

Description Display statistics about the Packet Forwarding Engine IP.

Options icmp—(Optional) Display Packet Forwarding Engine IP ICMP statistics.

options—(Optional) Display Packet Forwarding Engine IP options statistics.

sfm *slot*—(Optional) Show statistics for this SFM slot. *slot* can be a value from 0 through 3.

Required Privilege Level admin

Sample Output

```
user@host> show pfe statistics ip icmp
ICMP Statistics:
           0 requests
           0 network unreachables
           0 ttl expired
           0 ttl captured
           0 redirects
           0 mtu exceeded
           0 icmp/option handoffs

ICMP Errors:
           0 unknown unreachables
           0 unsupported ICMP type
           0 unprocessed redirects
           0 invalid ICMP type
           0 invalid protocol
           0 bad input interface
           0 throttled icmps
           0 runts

ICMP Discards:
           0 multicasts
           0 bad source addresses
           0 bad dest addresses
           0 IP fragments
           0 ICMP errors
User@host> show pfe statistics ip options
IP Option Values:
  LSRR/SSRR forwarding enabled

IP Option Statistics:
           0 loose source routes
           0 strict source routes
           0 record routes
      889382 router alerts
           0 other options

IP Option Errors:
           0 runts
           2 bad versions
           0 runt header lengths
           0 giant header lengths
           0 null frames
           0 bad option lengths
           0 duplicate options
           0 bad option pointers
           0 source route frames dropped
```

show pfe statistics notification

Syntax show pfe statistics notification

Description Display statistics about the Packet Forwarding Engine notification.

Required Privilege Level admin

Sample Output

```
user@host> show pfe statistics notification

PFE Notification statistics:
     919780 parsed
          0 aged
          0 corrupt
          0 illegal
          0 sample
          0 giants

PFE Notification Type statistics:
                Parsed       Input      Failed     Ignored
  Illegal           0           0           0           0
  Unclass       27636       27636           0           0
   Option      891032      890999          33           0
 Next-Hop        1112        1112           0           0
  Discard           0           0           0           0
   Sample           0           0           0           0
 Redirect           0           0           0           0
 DontFrag           0           0           0           0
```

show pfe statistics pio

Syntax show pfe statistics pio

Description Display statistics about the Packet Forwarding Engine polled I/O.

Required Privilege Level admin

Sample Output

```
user@host> show pfe statistics pio

PIO Statistics:

    8542732 PIO read requests
    8542732 PIO read replies
     586193 PIO write requests
     586191 PIO write replies
          0 PIO error replies
          0 PIO bad requests
          0 PIO bad replies
          0 PIO bad address
          0 PIO extra replies
          0 PIO timeouts
```

show pfe statistics traffic

Syntax show pfe statistics traffic

Description Display statistics about the Packet Forwarding Engine traffic.

Required Privilege Level admin

Output Fields PFE Traffic statistics—Information about PFE traffic.

- packets input—Number and rate of input packets.
- packets output—Number and rate of output packets.

PFE Local Traffic statistics—Information about PFE Local traffic.

- local packets input—Number of local input packets.
- local packets output—Number of local output packets.
- software input high drops—Number of software input high priority drops.
- software input medium drops—Number of software input medium priority drops.
- software input low drops—Number of software input low priority drops.
- software output drops—Number of software output drops.
- hardware input drops—Number of hardware input drops.

PFE Local Protocol statistics—Information about PFE Local Protocol.

- hdlc keepalives—Number of HDLC keepalive packets.
- atm oam—Number of ATM OAM packets.
- fr lmi—Number of Frame Relay LMI packets.
- ppp lcp/ncp—Number of PPP LCP/NCP packets.
- ospf hello—Number of OSPF hello packets.
- isis iih—Number of IS-IS IIH packets.

PFE Hardware Discard statistics—Infomation about PFE Hardware Discards.

- timeout—Number of packets discarded due to timeout.
- truncated key—Number of packets discarded due to truncated key.
- bits to test—Number of bits to test.
- data error—Number of packets discarded due to data error.
- stack underflow—Number of packets discarded due to stack underflow.
- stack overflow—Number of packets discarded due to stack overflow.

- discard route—Number of packets discarded due to discard route.
- illegal next hop—Number of packets discarded due to illegal next hop.
- invalid iif—Number of packets discarded due to invalid IIF.
- info cell drops—Number of info cell drops.
- fabric drops—Number of fabric drops.

Sample Output

```
user@host> show pfe statistics traffic
PFE Traffic statistics:
               910485 packets input  (34 packets/sec)
               932724 packets output (37 packets/sec)

PFE Local Traffic statistics:
     894563 local packets input
     932757 local packets output
          0 software input high drops
          0 software input medium drops
         25 software input low drops
          0 software output drops
          0 hardware input drops

PFE Local Protocol statistics:
       2578 hdlc keepalives
          0 atm oam
          0 fr lmi
       5191 ppp lcp/ncp
          0 ospf hello
         21 isis iih
PFE Hardware Discard statistics:
          0 timeout
          0 truncated key
          0 bits to test
          0 data error
          0 stack underflow
          0 stack overflow
       2638 discard route
          0 illegal next hop
```

show pfe terse

Syntax show pfe terse

Description Display information about the Packet Forwarding Engine.

Required Privilege Level admin

Sample Output

```
user@host> show pfe terse
Slot Type Slot    State   Flags Uptime
  0  SFM  Present Online  0x0bf 01:25:42
  2  SFM  Present Online  0x0bf 01:25:40
  0  FPC  Present Online  0x102 01:25:57
  1  FPC  Present Online  0x102 01:25:55
  2  FPC  Present Online  0x102 01:25:53
```

show system audit

Syntax show system audit <root-only>

Description Display the state and checksum values for all files in file systems. This command is equivalent to the UNIX mtree –Kmd5digest –c –p command.

To redirect the output to a file, issue the following command:

```
ssh router-name 'show system audit root-only' > output-file
```

If you save the output of the show system audit root-only command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

Options root-only—(Optional) Check only the root (/) file system. This command is equivalent to the UNIX mtree –Kmd5digest –c –p –x command.

Required Privilege Level admin

Sample Output

```
user@host> show system audit root-only
#          user: root
#       machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.               type=dir nlink=23 size=1024 time=950252640.0
    .cshrc      uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
                md5digest=f414e06fea6bd646244b98e13d6e6226
    .kernel.jkernel.backup \
                mode=0744 size=1934552 time=944688902.0 \
                md5digest=2c343cf0bd9fea8f04f78604feed7aa4
    .profile    uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
                md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
    COPYRIGHT   uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
                md5digest=7df8bc77dcee71382ea73eb0ec6a9243
    boot.config mode=0644 size=3 time=945902618.0 \
                md5digest=93d722493ed38477338a1405d7dcbb40
    boot.help   uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
                md5digest=9b7126385734bcae753f4179ab59d8e5
    compat      type=link mode=0777 size=11 time=915149058.0 \
                link=/usr/compat
    kernel      mode=0444 size=1947607 time=950230892.0 \
                md5digest=1a2a8aff2fec678a918ba0d6bf063980
    kernel.avr  uid=1112 size=1947642 time=950252597.0 \
                md5digest=82e1637682d58ec28964dfee7fccb62e
    kernel.config \
                mode=0644 size=0 time=915149058.0 \
                md5digest=d41d8cd98f00b204e9800998ecf8427e
    sys         type=link mode=0777 size=11 time=915149029.0 \
                link=usr/src/sys
```

show system boot-messages

Syntax show system boot-messages

Description Display initial messages generated by the system kernel upon boot. This is the contents of /var/run/dmesg.boot.

Required Privilege Level view

Sample Output

```
user@host> show system boot-messages
Copyright (c) 1992-1998 FreeBSD Inc.
Copyright (c) 1996-2000 Juniper Networks, Inc.
All rights reserved.
Copyright (c) 1982, 1986, 1989, 1991, 1993
        The Regents of the University of California.  All rights reserved.

JUNOS 4.1-20000216-Zf8469 #0: 2000-02-16 12:57:28 UTC
    tlim@single.juniper.net:/p/build/20000216-0905/4.1/release_kernel/sys/compil
e/GENERIC
CPU: Pentium Pro (332.55-MHz 686-class CPU)
  Origin = "GenuineIntel"  Id = 0x66a  Stepping=10
  Features=0x183f9ff<FPU,VME,DE,PSE,TSC,MSR,PAE,MCE,CX8,SEP,MTRR,PGE,MCA,CMOV,<b
16>,<b17>,MMX,<b24>>
Teknor CPU Card Recognized
real memory  = 805306368 (786432K bytes)
avail memory = 786280448 (767852K bytes)
Probing for devices on PCI bus 0:
chip0 <generic PCI bridge (vendor=8086 device=7192 subclass=0)> rev 3 class 6000
0 on pci0:0:0
chip1 <Intel 82371AB PCI-ISA bridge> rev 1 class 60100 on pci0:7:0
chip2 <Intel 82371AB IDE interface> rev 1 class 10180 on pci0:7:1
chip3 <Intel 82371AB USB interface> rev 1 class c0300 int d irq 11 on pci0:7:2
smb0 <Intel 82371AB SMB controller> rev 1 class 68000 on pci0:7:3
pcic0 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int a irq 15 on pci0:13
:0
TI1131 PCI Config Reg: [pci only][FUNC0 pci int]
pcic1 <TI PCI-1131 PCI-CardBus Bridge> rev 1 class 60700 int b irq 12 on pci0:13
:1
TI1131 PCI Config Reg: [pci only][FUNC1 pci int]
fxp0 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 12 on
 pci0:16:0
chip4 <generic PCI bridge (vendor=1011 device=0022 subclass=4)> rev 4 class 6040
0 on pci0:17:0
fxp1 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
 pci0:19:0
Probing for devices on PCI bus 1:
mcs0 <Miscellaneous Control Subsystem> rev 12 class ff0000 int a irq 12 on pci1:
13:0
fxp2 <Intel EtherExpress Pro 10/100B Ethernet> rev 8 class 20000 int a irq 10 on
 pci1:14:0
Probing for devices on the ISA bus:
sc0 at 0x60-0x6f irq 1 on motherboard
sc0: EGA color <16 virtual consoles, flags=0x0>
ed0 not found at 0x300
ed1 not found at 0x280
ed2 not found at 0x340
psm0 not found at 0x60
sio0 at 0x3f8-0x3ff irq 4 flags 0x20010 on isa
sio0: type 16550A, console
sio1 at 0x3e8-0x3ef irq 5 flags 0x20000 on isa
sio1: type 16550A
sio2 at 0x2f8-0x2ff irq 3 flags 0x20000 on isa
```

```
sio2: type 16550A
pcic0 at 0x3e0-0x3e1 on isa
PC-Card ctlr(0) TI PCI-1131 [CardBus bridge mode] (5 mem & 2 I/O windows)
pcic0: slot 0 controller I/O address 0x3e0
npx0 flags 0x1 on motherboard
npx0: INT 16 interface
fdc0: direction bit not set
fdc0: cmd 3 failed at out byte 1 of 3
fdc0 not found at 0x3f0
wdc0 at 0x1f0-0x1f7 irq 14 on isa
wdc0: unit 0 (wd0): <SunDisk SDCFB-80>, single-sector-i/o
wd0: 76MB (156672 sectors), 612 cyls, 8 heads, 32 S/T, 512 B/S
wdc0: unit 1 (wd1): <IBM-DCXA-210000>
wd1: 8063MB (16514064 sectors), 16383 cyls, 16 heads, 63 S/T, 512 B/S
wdc1 not found at 0x170
wdc2 not found at 0x180
ep0 not found at 0x300
fxp0: Ethernet address 00:a0:a5:12:05:5a
fxp1: Ethernet address 00:a0:a5:12:05:59
fxp2: Ethernet address 02:00:00:00:00:01
swapon: adding /dev/wd1s1b as swap device
Automatic reboot in progress...
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd0s1e: clean, 9233 free (9 frags, 1153 blocks, 0.1% fragmentation)
/dev/rwd0s1a: clean, 16599 free (95 frags, 2063 blocks, 0.1% fragmentation)
/dev/rwd1s1f: clean, 4301055 free (335 frags, 537590 blocks, 0.0% fragmentation)
```

show system buffers

Syntax show system buffers

Description Display information about the buffer pool that the Routing Engine uses for local traffic, which is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, pings, and so on).

Required Privilege Level view

Output Fields mbufs in use—Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. A special type of mbuf called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman. Each mbuf has a type, and the following lines itemize the amount allocated for each type. Types with no mbufs allocated are not displayed.

mbufs allocated to packet headers—Number of mbufs currently holding packet headers.

mbufs allocated to protocol control blocks—Number of mbufs currently holding state for sockets.

mbufs allocated to socket send data—Number of mbufs currently holding socket send data.

mbufs allocated to pfe refill data—Number of mbufs currently holding PFE refill data.

mbufs allocated to fxp data—Number of mbufs currently holding fxp data.

mbufs allocated to socket names and addresses—Number of mbufs currently holding addresses for sockets.

mbuf clusters in use—Allocation statistics for mbuf clusters.

allocated to network—Total amount of memory in use by the networking and IPC code.

requests for memory denied—Number of times a memory allocation request within the IPC and networking code failed.

requests for memory delayed—Number of times a memory allocation request within the IPC and networking code was postponed.

calls to protocol drain routines—Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

Sample Output

```
user@host> show system buffers
        853 mbufs in use:
        2 mbufs allocated to packet headers
        37 mbufs allocated to protocol control blocks
        28 mbufs allocated to socket names and addresses
        2 mbufs allocated to socket send data
        400 mbufs allocated to pfe refill data
        384 mbufs allocated to fxp data
784/944 mbuf clusters in use
1994 Kbytes allocated to network (83% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

show system connections

Syntax show system connections <extensive> <inet | inet6>

Description Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

Options none—Information about the active IP sockets on the Routing Engine. This option is equivalent to the UNIX netstat –an –f inet command.

extensive—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections. This option is equivalent to the UNIX netstat –Van –f inet command.

inet—(Optional) Display IPv4 connections.

inet6—(Optional) Display IPv6 connections.

Required Privilege Level view

Sample Output Sample Output: show system connections on page 90
Sample Output: show system connections extensive on page 90

Output Fields Proto—Protocol of the socket. It can be either TCP or UDP.

Recv-Q—Number of input packets received by the protocol and waiting to be processed by the application.

Send-Q—Number of output packets sent by the application and waiting to be processed by the protocol.

Local Address—Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.

Foreign Address—Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.

(state)—For TCP, the protocol state of the socket.

Sample Output: show system connections

```
user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp        0      2  192.168.4.16.513       208.197.169.254.894    ESTABLISHED
tcp        0      0  192.168.4.16.513       208.197.169.195.945    ESTABLISHED
tcp        0      0  *.23                   *.*                    LISTEN
tcp        0      0  *.22                   *.*                    LISTEN
tcp        0      0  *.513                  *.*                    LISTEN
tcp     0       0   *.514               *.*                    LISTEN
tcp     0       0   *.21                *.*                    LISTEN
tcp     0       0   *.79                *.*                    LISTEN
tcp     0       0   *.1023              *.*                    LISTEN
tcp     0       0   *.111               *.*                    LISTEN
udp     0       0   192.168.4.16.1634   208.197.169.249.2049
udp     0       0   192.168.4.16.1627   208.197.169.254.2049
udp     0       0   192.168.4.16.1371   208.197.169.195.2049
udp     0       0   *.*                 *.*
udp     0       0   *.9999              *.*
udp     0       0   *.161               *.*
udp     0       0   192.168.4.16.1039   192.168.4.16.1023
udp     0       0   192.168.4.16.1038   192.168.4.16.1023
udp     0       0   192.168.4.16.1037   192.168.4.16.1023
udp     0       0   192.168.4.16.1036   192.168.4.16.1023
udp     0       0   *.1022              *.*
udp     0       0   *.1023              *.*
udp     0       0   *.111               *.*
udp     0       0   *.*                 *.*
```

Sample Output: show system connections extensive

```
user@host> show system connections extensive
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp        0      2  192.168.4.16.513       208.197.169.254.894    ESTABLISHED
       iss: 3972677059      sndup: 3972693435       sndcc:         10
    snduna: 3972693435     sndnxt: 3972693437      sndwnd:      17376
    sndmax: 3972693437    sndcwnd:      65535 sndssthresh: 1073725440
       irs:  484187869      rcvup:  484188060       rcvcc:      98357
    rcvnxt:  484188070     rcvadv:  484205446      rcvwnd:      17376
       rtt:          1       srtt:          7        rttv:          5
    rttmin:          2   duration:       5011         mss:       1448
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP RCVD_CC [0x41e0]
tcp        0      0  192.168.4.16.513       208.197.169.195.945    ESTABLISHED
       iss: 1057609890      sndup: 1057790796       sndcc:          2
    snduna: 1057790810     sndnxt: 1057790810      sndwnd:      17376
    sndmax: 1057790810    sndcwnd:      39096 sndssthresh: 1073725440
       irs: 3551947312      rcvup: 3551947422       rcvcc:          0
    rcvnxt: 3551947422     rcvadv: 3551964798      rcvwnd:      17376
       rtt:          0       srtt:         17        rttv:         11
    rttmin:          2   duration:     125814         mss:       1448
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP [0x1e0]
udp     0      0    192.168.4.16.1634       208.197.169.249.2049
udp     0      0    192.168.4.16.1627       208.197.169.254.2049
udp     0      0    192.168.4.16.1371       208.197.169.195.2049
udp     0      0    *.*                     *.*
udp     0      0    *.9999                  *.*
udp     0      0    *.161                   *.*
udp     0      0    192.168.4.16.1039       192.168.4.16.1023
udp     0      0    192.168.4.16.1038       192.168.4.16.1023
udp     0      0    192.168.4.16.1037       192.168.4.16.1023
udp     0      0    192.168.4.16.1036       192.168.4.16.1023
udp     0      0    *.1022                  *.*
udp     0      0    *.1023                  *.*
udp     0      0    *.111                   *.*
udp     0      0    *.*                     *.*
```

show system processes

Syntax show system processes <wide> <brief | detail | extensive | summary>

Description Display information about software processes that are running on the router and that have controlling terminals. These commands are equivalent to various UNIX ps and top commands.

Options none—Display information about system processes. This option is equivalent to the UNIX ps –ax command.

brief—(Optional) Display brief system process information, listing no processes themselves. This option is equivalent to the UNIX top –bSd1 0 command.

detail—(Optional) Display detailed system process information. This option is equivalent to the UNIX ps –ax –rOuid,ppid,cpu,pri,nice,rss,wchan,start command.

extensive—(Optional) Display exhaustive system process information. This option is equivalent to the UNIX top –bSId1 infinity command.

summary—(Optional) Display a summary of active system process information. This option is equivalent to the beginning of the UNIX top command.

wide—(Optional) Display process information that might be wider than 80 columns. This option is equivalent to the UNIX ps –ax –ww command.

Required Privilege Level view

Sample Output Sample Output: show system processes (standard) on page 94
Sample Output: show system processes brief on page 95
Sample Output: show system processes detail on page 95
Sample Output: show system processes summary on page 95
Sample Output: show system processes extensive on page 96

Options at a Glance Table 7 summarizes which information is included in the output of each of the show system processes command options. In this table, output fields are listed in alphabetical order. In the Output Fields section, the fields are listed in the order in which they are displayed

Table 7: Show System Processes Output Field Summary

Options	Field Description
Standard Detail Summary Extensive	COMMAND—Command that is running.
Summary Extensive	CPU—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.
Detail	CPU—Short-term CPU usage.
Brief Summary Extensive	last pid—Last process identifier assigned to the process.
Brief Summary Extensive	load averages—Three load averages followed by the current time.
Brief Summary Extensive	Mem—Information about physical and virtual memory allocation.
Detail Summary Extensive	NI or Nice—UNIX “nice” value.
Standard Detail Summary Extensive	PID—Process identifier.
Detail	PPID—Parent process identifier.
Detail Summary Extensive	Pri—Current priority of the process. A lower number indicates a higher priority.
Brief Summary Extensive	processes—Number of existing processes and the number of processes in each state (sleeping, running, starting, zombies, and stopped).
Summary Extensive	RES—Current amount of resident memory in KB.
Detail	RSS—Resident set size.
Summary Extensive	SIZE—Total size of the process (text, data, and stack) in KB.
Detail	STARTED—Time when the process started running, given in local time.
Summary Extensive	STATE—Current state of the process (sleep, wait, run, idle, zombi, or stop).
Standard Detail	STAT—Symbolic process state. The state is given by a sequence of letters.
Brief Summary Extensive	Swap—Information about physical and virtual memory allocation.
Summary	TIME—Number of system and user CPU seconds that the process has used.
Standard Detail Extensive	TIME—Total amount of time that the command has been running.
Standard Detail	TT—Control terminal name.
Detail	UID—User identifier.
Summary Extensive	USERNAME—Name of the process’ owner.
Detail	WCHAN—Symbolic name of the wait channel.
Summary Extensive	WCPU—Weighted CPU usage.

Output Fields

last pid—Last process identifier assigned to the process.

load averages—Three load averages followed by the current time.

processes—Number of existing processes and the number of processes in each state (sleeping, running, starting, zombies, and stopped).

Mem—Information about physical and virtual memory allocation.

Swap—Information about physical and virtual memory allocation.

PID—Process identifier.

TT—Control terminal name.

STAT—Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process:

- D—In disk or other short-term, uninterruptible wait
- I—Idle (sleeping longer than about 20 seconds)
- R—Runnable
- S—Sleeping for less than 20 seconds
- T—Stopped
- Z—Dead (zombie)

Any additional characters indicate additional information:

- +—The process is in the foreground process group of its control terminal.
- <—The process has raised CPU scheduling priority.
- >—The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped.
- A—The process has asked for random page replacement.
- E—The process is trying to exit.
- L—The process has pages locked in core.
- N—The process has reduced CPU scheduling priority.
- S—The process has asked for FIFO page replacement.
- s—The process is a session leader.
- V—The process is suspended during a vfork.
- W—The process is swapped out.
- X—The process is being traced or debugged.

UID—User identifier.

USERNAME—Name of the process' owner.

PPID—Parent process identifier.

CPU—(Detail output only) Short-term CPU usage.

CPU—(Summary and Extensive output only) Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.

RSS—Resident set size.

WCHAN—Symbolic name of the wait channel.

STARTED—Time when the process started running, given in local time.

PRI—Current priority of the process. A lower number indicates a higher priority.

NI or NICE—UNIX "nice" value.

SIZE—Total size of the process (text, data, and stack) in KB.

RES—Current amount of resident memory in KB.

STATE—Current state of the process (sleep, wait, run, idle, zombi, or stop).

TIME—(Summary output only) Number of system and user CPU seconds that the process has used.

TIME—(Standard, Detail, and Extensive output only) Total amount of time that the command has been running.

WCPU—Weighted CPU usage.

COMMAND—Command that is currently running.

Sample Output: show system processes (standard)

```
user@host> show system processes
PID  TT  STAT      TIME COMMAND
    0  ??  DLs    0:00.70  (swapper)
    1  ??  Is     0:00.35 /sbin/init --
    2  ??  DL     0:00.00  (pagedaemon)
    3  ??  DL     0:00.00  (vmdaemon)
    4  ??  DL     0:42.37  (update)
    5  ??  DL     0:00.00  (if_jnx)
   80  ??  Ss     0:14.66 syslogd -s
   96  ??  Is     0:00.01 portmap
  128  ??  Is     0:02.70 cron
  173  ??  Is     0:02.24 /usr/local/sbin/sshd (sshd1)
  189  ??  S      0:03.80 /sbin/watchdog -t180
  190  ??  I      0:00.03 /usr/sbin/tnetd -N
  191  ??  S      2:24.76 /sbin/ifd -N
  192  ??  S<     0:55.44 /usr/sbin/xntpd -N
  195  ??  S      0:53.11 /usr/sbin/snmpd -N
  196  ??  S      1:15.73 /usr/sbin/mib2d -N
  198  ??  I      0:00.75 /usr/sbin/inetd -N
 2677  ??  I      0:00.01 /usr/sbin/mgd -N
 2712  ??  Ss     0:00.24 rlogind
```

```
 2735  ??  R      0:00.00 /bin/ps -ax
 1985  p0- S      0:07.41 ./rpd -N
 2713  p0  Is     0:00.24 -tcsh (tcsh)
 2726  p0  S+     0:00.07 cli
```

Sample Output: show system processes brief

```
user@host> show system processes brief
last pid:   543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

Sample Output: show system processes detail

```
user@host> show system processes detail
PID   UID  PPID CPU PRI NI  RSS WCHAN  STARTED   TT  STAT      TIME COMMAND
 3151  1049  3129   2  28  0  672 -       1:13PM   p0  R+     0:00.00 ps -ax -r
    1     0     0   0  10  0  376 wait    1:51PM   ??  Is     0:00.29 /sbin/ini
    2     0     0   0 -18  0   12 psleep  1:51PM   ??  DL     0:00.00  (pagedae
    3     0     0   0  28  0   12 psleep  1:51PM   ??  DL     0:00.00  (vmdaemo
    4     0     0   0  28  0   12 update  1:51PM   ??  DL     0:07.15  (update)
    5     0     0   0   2  0   12 pfesel  1:51PM   ??  IL     0:02.90  (if_pfe)
   27     0     1   0  10  0 17936 mfsidl  1:51PM   ??  Is     0:00.46 mfs /dev/
   81     0     1   0   2  0  496 select  1:52PM   ??  Ss     0:31.21 syslogd -
  119     1     1   0   2  0  492 select  1:52PM   ??  Is     0:00.00 portmap
  129     0     1   0  10  0  400 nfsidl  1:52PM   ??  I      0:02.15 nfsiod -n
  130     0     1   0  10  0  400 nfsidl  1:52PM   ??  I      0:02.04 nfsiod -n
  131     0     1   0  10  0  400 nfsidl  1:52PM   ??  I      0:00.01 nfsiod -n
  132     0     1  48  10  0  400 nfsidl  1:52PM   ??  I      0:00.00 nfsiod -n
  134     0     1   0   2  0  580 select  1:52PM   ??  S      0:02.95 amd -p -a
  151     0     1   0  18  0  532 pause   1:52PM   ??  Is     0:00.34 cron
  183     0     1   0   2  0  420 select  1:52PM   ??  Ss     0:00.07 /usr/loca
  206     0     1   0  18  0   72 pause   1:52PM   ??  S      0:00.51 /sbin/wat
  207     0     1   0   2  0  520 select  1:52PM   ??  I      0:00.16 /usr/sbin
  208     0     1   0   2  0  536 select  1:52PM   ??  S      0:08.21 /sbin/dcd
  210     0     1 255   2 -12  740 select  1:52PM   ??  S<     0:05.83 /usr/sbin
  211     0     1   0   2  0  376 select  1:52PM   ??  S      0:00.03 /usr/sbin
  215     0     1   0   2  0  548 select  1:52PM   ??  I      0:00.50 /usr/sbin
  219     0     1   0   3  0  540 ttyin   1:52PM   v0  Is+    0:00.02 /usr/libe
  220     0     1   0   3  0  540 ttyin   1:52PM   v1  Is+    0:00.01 /usr/libe
  221     0     1   0   3  0  540 ttyin   1:52PM   v2  Is+    0:00.01 /usr/libe
  222     0     1   0   3  0  540 ttyin   1:52PM   v3  Is+    0:00.01 /usr/libe
  735     0     1   0   2  0  468 select  2:47PM   ??  S      0:19.14 /usr/sbin
  736     0     1   0   2  0  212 select  2:47PM   ??  S      0:14.13 /usr/sbin
 1380     0     1   0   3  0  888 ttyin   7:32PM   d0  Is+    0:00.46 bash
 3019     0   207   0   2  0  636 select 10:49AM   ??  Ss     0:02.93 tnp.chass
 3122     0  1380   0   2  0 1764 select 12:33PM   d0  S      0:00.77 ./rpd -N
 3128     0   215   0   2  0  580 select 12:45PM   ??  Ss     0:00.12 rlogind
 3129  1049  3128   0  18  0  944 pause  12:45PM   p0  Ss     0:00.14 -tcsh (tc
    0     0     0   0 -18  0    0 sched   1:51PM   ??  DLs    0:00.10  (swapper
```

Sample Output: show system processes summary

```
user@host> show system processes summary
last pid:   543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  527 root      2   0   176K   580K select   0:00  0.04%  0.04% rlogind
  543 root     30   0   604K   768K RUN      0:00  0.00%  0.00% top
```

Sample Output: show system processes extensive

```
user@host> show system processes extensive
last pid:   544;  load averages:  0.00,  0.00,  0.00    18:30:33
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  544 root     30   0   604K   768K RUN      0:00  0.00%  0.00% top
    3 root     28   0     0K    12K psleep   0:00  0.00%  0.00% vmdaemon
    4 root     28   0     0K    12K update   0:03  0.00%  0.00% update
  528 aviva    18   0   660K   948K pause    0:00  0.00%  0.00% tcsh
  204 root     18   0   300K   544K pause    0:00  0.00%  0.00% csh
  131 root     18   0   332K   532K pause    0:00  0.00%  0.00% cron
  186 root     18   0   196K    68K pause    0:00  0.00%  0.00% watchdog
   27 root     10   0   512M 16288K mfsidl   0:00  0.00%  0.00% mount_mfs
  112 root     10   0   148K   400K nfsidl   0:00  0.00%  0.00% nfsiod
  110 root     10   0   148K   400K nfsidl   0:00  0.00%  0.00% nfsiod
  111 root     10   0   148K   400K nfsidl   0:00  0.00%  0.00% nfsiod
  109 root     10   0   148K   400K nfsidl   0:00  0.00%  0.00% nfsiod
    1 root     10   0   620K   344K wait     0:00  0.00%  0.00% init
  304 root      3   0   884K   900K ttyin    0:00  0.00%  0.00% bash
  200 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  203 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  202 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  201 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  194 root      2   0  2248K  1640K select   0:11  0.00%  0.00% rpd
  205 root      2   0   964K   800K select   0:12  0.00%  0.00% tnp.chassisd
  189 root      2 -12   352K   740K select   0:03  0.00%  0.00% xntpd
  114 root      2   0   296K   612K select   0:00  0.00%  0.00% amd
  188 root      2   0   780K   600K select   0:00  0.00%  0.00% dcd
  527 root      2   0   176K   580K select   0:00  0.00%  0.00% rlogind
  195 root      2   0   212K   552K select   0:00  0.00%  0.00% inetd
  187 root      2   0   192K   532K select   0:00  0.00%  0.00% tnetd
   83 root      2   0   188K   520K select   0:00  0.00%  0.00% syslogd
  538 root      2   0  1324K   516K select   0:00  0.00%  0.00% mgd
   99 daemon    2   0   176K   492K select   0:00  0.00%  0.00% portmap
  163 root      2   0   572K   420K select   0:00  0.00%  0.00% nsrexecd
  192 root      2   0   560K   400K select   0:10  0.00%  0.00% snmpd
  191 root      2   0  1284K   376K select   0:00  0.00%  0.00% mgd
  537 aviva     2   0   636K   364K select   0:00  0.00%  0.00% cli
  193 root      2   0   312K   204K select   0:07  0.00%  0.00% mib2d
    5 root      2   0     0K    12K pfesel   0:00  0.00%  0.00% if_pfe
    2 root    -18   0     0K    12K psleep   0:00  0.00%  0.00% pagedaemon
    0 root    -18   0     0K     0K sched    0:00  0.00%  0.00% swapper
```

show system reboot

Syntax show system reboot

Description Display any pending system reboots or halts.

Required Privilege Level maintenance

Sample Output

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

show system software

Syntax show system software <detail>

Description Display loaded JUNOS extensions.

Options detail—(Optional) Display detailed information about available JUNOS extensions.

Required Privilege Level maintenance

show system statistics

Syntax show system statistics *<protocol>*

Description Display systemwide protocol-related statistics.

Options *protocol*—(Optional) Name of a protocol:

- arp—Address Resolution Protocol
- clnl—Connectionless Network Layer
- esis—End System–Intermediate System
- icmp—Internet Control Message Protocol
- icmpv6—Internet Control Message Protocol version 6
- igmp—Internet Group Management Protocol
- ip—Internet Protocol version 4
- ipv6—Internet Protocol version 6
- mpls—Multiprotocol Label Switching
- rdp—Reliable Datagram Protocol
- tcp—Transmission Control Protocol
- tnp—Trivial Network Protocol
- tudp—Trivial User Datagram Protocol
- udp—User Datagram Protocol

Required Privilege Level view

Sample Output

```
user@host> show system statistics
ip:
        476559 total packets received
        0 bad header checksums
        0 with size smaller than minimum
        0 with data size < data length
        0 with header length < data size
        0 with data length < header length
        0 with bad options
        0 with incorrect version number
        0 fragments received
        0 fragments dropped (dup or out of space)
        0 fragments dropped after timeout
        0 packets reassembled ok
        406456 packets for this host
        39565 packets for unknown/unsupported protocol
        354 packets forwarded
        0 packets not forwardable
        0 redirects sent
        342678 packets sent from this host
        5 packets sent with fabricated ip header
        0 output packets dropped due to no bufs, etc.
        0 output packets discarded due to no route
        0 output datagrams fragmented
        0 fragments created
        0 datagrams that can't be fragmented
icmp: ...
tcp: ...
udp: ...
igmp: ...
arp: ...
clnl: ...
esis: ...
```

show system storage

Syntax show system storage

Description Display statistics about the amount of free disk space in the router's file systems. Values are displayed in 1024-byte (1-KB) blocks. This command is equivalent to the UNIX df –k command.

Required Privilege Level view

Sample Output

```
user@host> show system storage
Filesystem                                1K-blocks     Used    Avail Capacity  Mon
/dev/wd1s1a                                   65687    28888    31545    48%    /
devfs                                            16       16        0   100%    dut
devfs                                            16       16        0   100%    /d/
/dev/vn1                                       9352     9352        0   100%    /pe
/dev/vn2                                       1442     1442        0   100%    /pr
/dev/vn3                                        694      694        0   100%    /pr
/dev/vn4                                       8094     8094        0   100%    /pr
/dev/vn5                                      11194    11194        0   100%    /pr
/dev/vn6                                       5652     5652        0   100%    /pr
mfs:1577                                    1015815        2   934548     0%    /tp
/dev/wd1s1e                                   25263       14    23228     0%    /cg
procfs                                            4        4        0   100%    /pc
/dev/wd1s1f                                 9811371    36499  8989963     0%    /vr
pid1836@sage:/home                                0        0        0   100%    /he
pid1836@sage:/host                                0        0        0   100%    /ht
pid1836@sage:/volume                              0        0        0   100%    /ve
pid1836@sage:/homes                               0        0        0   100%    /hs
green:/vol/vol0/freebsd-bin/X11R6_4      10485760  4669292  5816468    45%    /.4
yellow:/vol/vol0/freebsd-bin/local_4x    10485760  1191836  9293924    11%    /.x
blue:/vol/vol0/homes                    324492812296203268 28289544    91%    /s
red:/usr/home                             7546483  5012946  1929819    72%    /.e
green:/vol/vol0/labtools                 10485760  3331396  7154364    32%    /.s
green:/vol/vol0/buildtools                6291456     6444  6285012     0%    /.s
```

Output Fields

Filesystem—Name of the file system. wd0 is the flash drive, wd2 is the hard drive, and wfd0 is the LS-120 floppy drive.

1K-blocks—Size of the file system in 1-KB blocks.

Used—Amount of space used in the file system, in blocks.

Avail—Amount of space available in the file system, in blocks.

Capacity—Percentage of the file system's space that is being used.

Mounted on—Directory to which the file system is mounted.

show system uptime

Syntax show system uptime

Description Display the current time and information about how long the router, router software, and routing protocols have been running.

Required Privilege Level view

Output Fields Current time—Current system time in UTC.

System booted—Date and time when the router was last booted and how long it has been running.

Protocols started—Date and time when the routing protocols were last started and how long they have been running.

Last configured—Date and time when a configuration was last committed. Also shows name of user who issued the last commit command.

time and up—Current time, in the local time zone, and how long the router has been operational.

users—Number of users logged into the router.

load averages—Load averages for the last 1 minute, 5 minutes, and 15 minutes.

Sample Output

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM  up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system users

Syntax show system users <noresolve>

Description List information about the users who are currently logged into the router.

Options noresolve—(Optional) Do not attempt to resolve IP addresses to host names.

Required Privilege Level view

Output Fields *time* and up—Current time, in the local time zone, and how long the router has been operational.

users—Number of users logged into the router.

load averages—Load averages for the last 1 minute, 5 minutes, and 15 minutes.

USER—Username.

TTY—Terminal through which the user is logged in.

FROM—System from which the user has logged in. A hyphen indicates that the user is logged in through the console.

LOGIN@—Time when the user logged in.

IDLE—How long the user has been idle.

WHAT—Processes that the user is running.

Sample Output

```
user@host> show system users
 7:30PM  up 4 days,  2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER     TTY FROM              LOGIN@  IDLE WHAT
root     d0  -                Fri05PM 4days -csh (csh)
blue   p0 level5.company.net 7:30PM     - cli
```

show system virtual-memory

Syntax show system virtual-memory

Description Display the usage of JUNOS kernel memory listed first by size of allocation and then by type of usage. show system virtual-memory is to be used for troubleshooting with JTAC.

Privilege Level view

Output Fields Memory statistics by bucket size:

- Size—Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
- In Use—Number of memory blocks of this size that are in use (bytes).
- Free—Number of memory blocks of this size that are free (bytes).
- Requests—Number of memory allocation requests made so far.
- HighWater—The maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
- Couldfree—The total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.

Memory usage type by bucket size:

- Size—Memory block size (bytes).
- Type(s)—Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.

Memory statistics by type:

- Type—Kernel module that is using dynamic memory.
- InUse—Number of memory blocks used by this type. The number is rounded up.
- HighUse—Maximum memory ever used by this type.
- Limit—Maximum memory that can be allocated to this type.
- Requests—Total number of dynamic memory allocation requests this type has made.
- Type Limit—Number of times requests were blocked for reaching the maximum limit.
- Kern Limit—Number of times requests were blocked for kernel map.
- Size(s)— Memory block sizes this type is using.

Memory Totals:

- In Use—Total kernel dynamic memory in use (bytes, rounded up).
- Free—Total kernel dynamic memory free (bytes, rounded up).
- Requests—Total number of memory allocation requests.

Sample Output

```
user@host> show system virtual-memory
Memory statistics by bucket size
Size   In Use   Free   Requests  HighWater  Couldfree
  16      658    110     137296       0       1280
  32      423    217     637517       0        640
  64     1783   1033     964158       0        320
 128     1412     92     333253       0        160
 256      851     29     192616       0         80
 512      120      8       3524       0         40
  1K      103      5      12174       0         20
  2K       66     70       6797       0         10
  4K       34      0        207       0          5
  8K        3      1       7183       0          5
 16K        6      0          6       0          5
 32K        2      0       4310       0          5
 64K       20      0         24       0          5
128K        5      0         15       0          5
256K        2      0          2       0          5
512K        1      0          1       0          5

Memory usage type by bucket size
Size  Type(s)
  16  uc_devlist, COS, BPF, DEVFS mount, DEVFS node, vnodes, mount, pcb,
          soname, rman, bus, sysctl, pfe_ipc, mkey, ifmaddr, ipfw,
          rnode, temp,
          devbuf, atexit, proc-args, kld, MD disk
  32  atkbddev, Gzip trees, dirrem, mkdir, diradd, freefile, freefrag,
          indirdep, bmsafemap, newblk, tseg_qent, COS, vnodes,
          cluster_save buffer, pcb, soname, taskqueue, SWAP,
          eventhandler, bus,
          sysctl, uidinfo, subproc, pgrp, ifstate, pfe_ipc, mkey, ifmaddr,
          rnode, temp, devbuf, proc-args, sigio, kld
  64  isadev, AD driver, Gzip trees, MFS node, allocindir, allocdirect,
          pagedep, NFS req, DEVFS name, vnodes, cluster_save buffer,
          vfscache,
          pcb, rman, eventhandler, bus, sysctl, subproc, session, pfe_ipc,
          mkey, ifaddr, rtable, rtnexthop, iffamily, temp, devbuf, lockf,
          proc-args, file
 128  ZONE, Gzip trees, freeblks, inodedep, COS, DEVFS node, vnodes, mount,
          cluster_save buffer, vfscache, pcb, soname, ttys, bus,
          uidinfo, cred,
          ifstate, pfe_ipc, ifstat, metrics, rtnexthop, temp, devbuf,
          zombie,
          proc-args, dev_t, timecounter, kld
 256  Gzip trees, MFS node, FFS node, newblk, NFS daemon, vnodes, pcb,
          ISOFS node, ttys, bus, subproc, pfe_ipc, sysctl, rtable, ipfw,
          iflogical, iftable, temp, devbuf, proc-args, file desc
 512  ATA generic, Gzip trees, UFS mount, NFSV3 diroff, mount, BIO buffer,
          pcb, ptys, msg, ioctlops, bus, proc, ISOFS mount, pfe_ipc, tagnh,
          temp, devbuf, file desc, MD disk
  1K  AD driver, Gzip trees, NQNFS Lease, BIO buffer, sem, bus, pfe_ipc,
          ifstat, temp, devbuf, file desc, kld
  2K  uc_devlist, Gzip trees, UFS mount, BIO buffer, pcb, bus, rcache,
          temp, devbuf, file desc, kld
```

```
  4K  memdesc, Gzip trees, sem, msg, bus, temp, devbuf, kld
  8K  Gzip trees, UFS mount, indirdep, shm, bus, devbuf
 16K  msg, temp, devbuf
 32K  Gzip trees, pagedep, devbuf
 64K  mbuf, MSDOSFS mount, devbuf
128K  UFS ihash, inodedep, NFS hash, bus, ISOFS mount, kld
256K  VM pgdata, vfscache
512K  SWAP

Memory statistics by type                          Type  Kern
        Type  InUse MemUse HighUse  Limit Requests Typlm Krnlm Size(s)
       isadev    16     1K      1K102400K       16    0     0  64
     atkbddev     2     1K      1K102400K        2    0     0  32
   uc_devlist    24     3K      3K102400K       24    0     0  16,2K
      memdesc     1     4K      4K102400K        1    0     0  4K
         mbuf     1    40K     40K102400K        1    0     0  64K
    AD driver     2     2K      3K102400K    18185    0     0  64,1K
  ATA generic     1     1K      1K102400K        2    0     0  512
         ZONE    13     2K      2K102400K       13    0     0  128
    VM pgdata     1   256K    256K102400K        1    0     0  256K
   Gzip trees     0     0K     46K102400K   353636    0     0
32,64,128,256,512,1K,2K,4K,8K,32K
    UFS mount    12    24K     24K102400K       12    0     0  512,2K,8K
    UFS ihash     1   128K    128K102400K        1    0     0  128K
     MFS node     4     1K      3K102400K     9653    0     0  64,256
     FFS node   408   102K    115K102400K     5359    0     0  256
       dirrem     0     0K      3K102400K     1342    0     0  32
        mkdir     0     0K      1K102400K      104    0     0  32
       diradd     0     0K      3K102400K     1336    0     0  32
     freefile     0     0K      3K102400K      695    0     0  32
     freeblks     0     0K      7K102400K      799    0     0  128
     freefrag     0     0K      2K102400K     1052    0     0  32
   allocindir     0     0K     70K102400K     6150    0     0  64
     indirdep     0     0K     17K102400K      183    0     0  32,8K
  allocdirect     1     1K     10K102400K     2824    0     0  64
    bmsafemap     1     1K      1K102400K      973    0     0  32
       newblk     1     1K      1K102400K     8975    0     0  32,256
     inodedep     2   129K    158K102400K     2918    0     0  128,128K
      pagedep     1    32K     33K102400K      701    0     0  64,32K
     NFS hash     1   128K    128K102400K        1    0     0  128K
  NQNFS Lease     1     1K      1K102400K        1    0     0  1K
 NFSV3 diroff     3     2K      2K102400K       10    0     0  512
   NFS daemon     1     1K      1K102400K        1    0     0  256
      NFS req     0     0K      1K102400K    13560    0     0  64
    tseg_qent     0     0K      1K102400K      112    0     0  32
          COS    27     2K      2K102400K      167    0     0  16,32,128
          BPF   101     2K      2K102400K      101    0     0  16
MSDOSFS mount     1    64K     64K102400K        1    0     0  64K
  DEVFS mount     2     1K      1K102400K        2    0     0  16
   DEVFS name   471    30K     30K102400K      471    0     0  64
   DEVFS node   457    56K     56K102400K      466    0     0  16,128
       vnodes    85     9K      9K102400K      736    0     0
16,32,64,128,256
        mount    27    14K     15K102400K      101    0     0  16,128,512
cluster_save buffer     0     0K      1K102400K      149    0     0
32,64,128
     vfscache   586   294K    295K102400K    21142    0     0  64,128,256K
   BIO buffer   106   138K    274K102400K     4457    0     0  512,1K,2K
          pcb   188    36K     39K102400K    13808    0     0
```

```
16,32,64,128,256,512,2K
       soname    72     6K      6K102400K   215980    0     0  16,32,128
   ISOFS node   234    59K     63K102400K     1564    0     0  256
         ptys    32    16K     16K102400K       32    0     0  512
         ttys   369    47K     47K102400K      798    0     0  128,256
          shm     1     8K      8K102400K        1    0     0  8K
          sem     3     6K      6K102400K        3    0     0  1K,4K
          msg     4    25K     25K102400K        4    0     0  512,4K,16K
         rman    45     3K      3K102400K      115    0     0  16,64
     ioctlops     0     0K      1K102400K        4    0     0  512
    taskqueue     1     1K      1K102400K        1    0     0  32
         SWAP     2   413K    413K102400K        2    0     0  32,512K
 eventhandler    14     1K      1K102400K       14    0     0  32,64
          bus   381    26K    109K102400K      644    0     0
16,32,64,128,256,512,1K,2K,4K,8K,128K
       sysctl     0     0K      1K102400K    24431    0     0  16,32,64
      uidinfo     5     1K      1K102400K       23    0     0  32,128
         cred    25     4K      4K102400K   149767    0     0  128
      subproc   139    11K     11K102400K   317428    0     0  32,64,256
         proc     2     1K      1K102400K        2    0     0  512
      session    21     2K      2K102400K     1782    0     0  64
         pgrp    26     1K      1K102400K     1847    0     0  32
  ISOFS mount     9   132K    133K102400K       11    0     0  512,128K
      ifstate    35     5K      5K102400K      148    0     0  32,128
      pfe_ipc     0     0K      1K102400K     4256    0     0
16,32,64,128,256,512,1K
         mkey   166     4K      4K102400K    11286    0     0  16,32,64
       ifaddr     9     1K      1K102400K       37    0     0  64
       sysctl     0     0K      1K102400K       12    0     0  256
       rtable    37     3K      3K102400K      124    0     0  64,256
      ifmaddr   108     3K      3K102400K      136    0     0  16,32
         ipfw     1     1K      1K102400K        2    0     0  16,256
       ifstat    32    12K     12K102400K      154    0     0  128,1K
        tagnh     1     1K      1K102400K        1    0     0  512
       rcache     3     6K      6K102400K        3    0     0  2K
        rnode    24     1K      1K102400K       90    0     0  16,32
      metrics     6     1K      1K102400K       14    0     0  128
    rtnexthop    36     4K      4K102400K       92    0     0  64,128
     iffamily     8     1K      1K102400K       22    0     0  64
    iflogical     5     2K      2K102400K       19    0     0  256
      iftable    13     4K      4K102400K       20    0     0  256
         temp   259    39K     40K102400K   136003    0     0
16,32,64,128,256,512,1K,2K,4K,16K
       devbuf   129  1410K   1434K102400K      522    0     0
16,32,64,128,256,512,1K,2K,4K,8K,16K,32K,64K
        lockf    22     2K      2K102400K    10038    0     0  64
       atexit     1     1K      1K102400K        1    0     0  16
       zombie     0     0K      1K102400K   158615    0     0  128
    proc-args    50     2K      3K102400K   131147    0     0
16,32,64,128,256
        sigio     1     1K      1K102400K       98    0     0  32
         file   245    16K     16K102400K   502481    0     0  64
    file desc    66    18K     26K102400K   158755    0     0  256,512,1K,2K
        dev_t   277    35K     35K102400K      277    0     0  128
  timecounter    10     2K      2K102400K       10    0     0  128
          kld     9    95K    100K102400K       21    0     0
16,32,128,1K,2K,4K,128K
      MD disk     2     1K      1K102400K        2    0     0  16,512

Memory Totals:  In Use    Free    Requests
                 3908K    249K     2299083
```

show task

Syntax show task <*task-name*> <summary> <io | memory <*object*>>

Description List the routing, routing protocol, and interface tasks that are currently running.

Options io—(Optional) Display task I/O statistics.

memory <*object*>—(Optional) Display detailed information about the memory utilization of running tasks or about the memory utilization of a particular task.

summary—(Optional) Display summary information about the running tasks.

task-name—(Optional) Display detailed information about a particular running task.

Default: summary

Required Privilege Level view

Sample Output Sample Output: show task memory on page 107
Sample Output: show task summary on page 108

Output Fields Pri—Current priority of the process. A lower number indicates a higher priority.

Task Name—Name of the task.

Pro—IP protocol number associated with the process.

Port—TCP or UDP port number associated with the task.

So—Socket number of the task.

Flags—Flags for the task:

- Accept—Task is waiting for incoming connections.
- Connect—Task is waiting for a connection to complete.
- Delete—Task has been deleted and is being cleaned up.
- LowPrio—Task will be dispatched to read its socket after other, higher priority tasks.

Sample Output: show task memory

```
user@host> show task memory
Memory                 Size (kB)  %Available  When
  Currently In Use:        29417          3%  now
  Maximum Ever Used:       33882          4%  00/02/11 22:07:03
  Available:              756281        100%  now
```

Sample Output: show task summary

```
user@host> show task summary
Pri Task Name                          Pro  Port So Flags
 10 IF
 15 LABEL
 15 ISO
 15 INET                                            7
 20 Aggregate
 20 RT
 30 ICMP                                  1         9
 39 ISIS I/O                                       12
 40 IS-IS                                          10
 40 BGP RT Background                                 <LowPrio>
 40 BGP.0.0.0.0+179                           179 15 <Accept LowPrio>
 50 BGP_69.192.168.201.234+179                179 17 <LowPrio>
 50 BGP_70.192.168.201.233+179                179 16 <LowPrio>
 50 BGP_Group_69_153                                  <LowPrio>
 50 BGP_Group_70_153                                  <LowPrio>
 50 ASPaths
 60 KRT                                   255       1
 60 Redirect
 70 MGMT.local                                     14 <LowPrio>
 70 MGMT_Listen./var/run/rpd_mgmt                  13 <Accept LowPrio>
 70 SNMP Subagent./var/run/sub_rpd.sock             8 <LowPrio>
```

show version

Syntax show version <brief | detail>

Description Display the host name and version information about the software running on the router.

Options none—Display version information for the JUNOS software packages and the software for each software process.

brief—(Optional) Display version information for the JUNOS software packages (that is, for the JUNOS, kernel, and Packet Forwarding Engine software).

detail—(Optional) Display software version information for JUNOS and other software packages running on the router.

Required Privilege Level view

Sample Output: show version

```
user@host> show version
Hostname: router1
Model: m10
JUNOS Base OS boot [5.3R1.2]
JUNOS Base OS Software Suite [5.3R1.2]
JUNOS Kernel Software Suite [5.3R1.2]
JUNOS Packet Forwarding Engine Support [5.3R1.2]
JUNOS Routing Software Suite [5.3R1.2]
JUNOS Online Documentation [5.3R1.2]
JUNOS Crypto Software Suite [5.3R1.2]
KERNEL 5.3R1.2 #0 built by builder on 2002-04-30 01:40:52 UTC
MGD release 5.3R1.2 built by builder on 2002-04-30 01:22:23 UTC
CLI release 5.3R1.2 built by builder on 2002-04-30 01:20:54 UTC
CHASSISD release 5.3R1.2 built by builder on 2002-04-30 01:20:11 UTC
DCD release 5.3R1.2 built by builder on 2002-04-30 01:19:26 UTC
RPD release 5.3R1.2 built by builder on 2002-04-30 01:23:18 UTC
SNMPD release 5.3R1.2 built by builder on 2002-04-30 01:28:00 UTC
MIB2D release 5.3R1.2 built by builder on 2002-04-30 01:22:27 UTC
APSD release 5.3R1.2 built by builder on 2002-04-30 01:20:07 UTC
VRRPD release 5.3R1.2 built by builder on 2002-04-30 01:28:12 UTC
ALARMD release 5.3R1.2 built by builder on 2002-04-30 01:20:01 UTC
PFED release 5.3R1.2 built by builder on 2002-04-30 01:23:03 UTC
CRAFTD release 5.3R1.2 built by builder on 2002-04-30 01:21:06 UTC
SAMPLED release 5.3R1.2 built by builder on 2002-04-30 01:27:48 UTC
ILMID release 5.3R1.2 built by builder on 2002-04-30 01:21:57 UTC
BPRELAYD release 5.3R1.2 built by builder on 2002-04-30 01:19:18 UTC
RMOPD release 5.3R1.2 built by builder on 2002-04-30 01:23:10 UTC
COSD release 5.3R1.2 built by builder on 2002-04-30 01:21:01 UTC
KMD release 5.3R1.2 built by builder on 2002-04-30 01:22:09 UTC
FSAD release 5.3R1.2 built by builder on 2002-04-30 01:21:23 UTC
SERVICED release 5.3R1.2 built by builder on 2002-04-30 01:27:53 UTC
IRSD release 5.3R1.2 built by builder on 2002-04-30 01:22:02 UTC
NASD release 5.3R1.2 built by builder on 2002-04-30 01:22:59 UTC
FUD release 5.3R1.2 built by builder on 2002-04-30 01:21:26 UTC
PPMD release 5.3R1.2 built by builder on 2002-04-30 01:23:08 UTC
jkernel-dd release 5.3R1.2 built by builder on 2002-04-30 01:16:30 UTC
jroute-dd release 5.3R1.2 built by builder on 2002-04-30 01:16:44 UTC
jdocs-dd release 5.3R1.2 built by builder on 2002-04-30 01:14:58 UTC
jcrypto-dd release 5.3R1.2 built by builder on 2002-04-30 01:16:58 UTC
```

ssh

Syntax

ssh *host* < bypass-routing > <routing-instance *routing-instance-name*> <source *address*> < vpn-interface *vpn-interface*> <v2>

Description

Use the secure shell (SSH) program to open a connection between a local router and a remote system and execute commands on the remote system. You can issue the ssh command from the JUNOS CLI to log into a remote system or from a remote system to log into the local router.

To open an SSH connection from a remote system to the router, use the ssh command on the remote system. When executing this command, you include one or more CLI commands by including them in quotation marks and separating the commands with semicolons:

ssh *address* '*cli-command1*; *cli-command2*'

To configure an SSH (version 1) key for your user account, include the authentication ssh-rsa statement at the [edit system login user *user-name*] hierarchy level, as described in the *JUNOS Internet Software Configuration Guide: Getting Started*.

To configure an SSH (version 2) key for your user account, include the authentication dsa-rsa statement at the [edit system login user *user-name*] hierarchy level, as described in the *JUNOS Internet Software Configuration Guide: Getting Started*.

If you use ssh to connect between a router running JUNOS Release 5.*x* and one running Release 4.*x*, the console displays warnings because of minor discrepancies between the two implementations: "Warning: Server lies about size of server host key: actual size is 1023 bits vs. announced 1024. Warning: This may be due to an old implementation of ssh." The warnings are informational and do not affect ssh operation.

Options

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

host—Name or address of the remote system.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the ssh attempt.

source *address*—(Optional) Source address of the ssh connection.

v2—(Optional) Use ssh version 2 only when connecting to a remote host.

vpn-interface *vpn-interface*—(Optional) VPN interface for SSH session. This option does not work when default-address-selection is configured.

Required Privilege Level

network

Sample Output

```
user@host> ssh berry
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes
Host 'berry' added to the list of known hosts.
boojun@berry's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
% exit
user@host>
```

start

Syntax start shell (csh | sh)

Description Exit from the CLI environment and create a UNIX-level shell.

To return to the CLI, type exit from the shell.

When you are in the shell, the shell prompt has the following format. Note that there is a space after the %.

username@hostname%

An example of the prompt is:

root@router%

Options csh—Create a UNIX C shell.

sh—Create a UNIX Bourne shell.

Required Privilege Level shell and maintenance

Sample Output

```
user@host> start shell csh
%
% exit
username@hostname% start shell sh
%
% exit
user@host>
```

telnet

Syntax telnet *host* <8bit> < bypass-routing > <noresolve> <port *port*> <routing-instance *routing-instance-name*> <vpn-interface *vpn-interface*> <source *address*>

Description Open a telnet session to a remote system.

Type Ctrl-] to escape from the telnet session to the telnet command level, and then type quit to exit from telnet.

Options 8bit—(Optional) Use an 8-bit data path.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

host—Name or address of the remote system.

noresolve—(Optional) Do not attempt to determine the host name that corresponds to the IP address.

port *port*—(Optional) Port number or service name on remote system.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the telnet attempt.

source *source*—(Optional) Source address of the telnet connection.

vpn-interface *vpn-interface*—(Optional) VPN interface for SSH session. This option does not work when default-address-selection is configured.

Required Privilege Level network

Sample Output

```
user@host> telnet 192.154.1.254
Trying 192.154.169.254...
Connected to level5.company.net.
Escape character is '^]'.
ttypa
login: <Control>]
telnet> quit
Connection closed.
user@host>
```

test configuration

Syntax test configuration (*filename* | terminal)

Description Verify that the syntax of a configuration file is correct.

If the configuration contains any errors, a message is displayed to indicate the line number and column number in which the error was found.

Options *filename*—Name of the configuration file. Specify the filename as described in "How to Specify Filenames and URLs" on page 43.

terminal—Use the text you type at the terminal as input to the configuration. Type Ctrl-D to end terminal input.

Required Privilege Level view

Sample Output

```
user@host> test configuration terminal
[Type ^D to end input]
system {
    host-name bluesky;
    paris-23;
    login;
}
terminal:3:(8) syntax error: paris
  [edit system]
    'paris-23;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
  [edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed
```

traceroute

Syntax traceroute *host* < as-number-lookup > < bypass-routing > <gateway *address*> < inet > < inet6 > <noresolve> <routing-instance *routing-instance-name*> <source *address*> <tos *value*> <ttl *value*> <vpn-interface *vpn-interface*> <wait *seconds*>

Description Display the route packets take to a specified network host. traceroute can be as a debugging tool to locate points of failure in a network.

Options as-number-lookup—(Optional) Display the autonomous system number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

host—IP address or name of remote host.

gateway *address*—(Optional) Address of a router through which the route should transit.

inet—(Optional) Trace the route belonging to IPv4.

inet6—(Optional) Trace the route belonging to IPv6.

noresolve—(Optional) Do not attempt to determine the host name that corresponds to the IP address.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type of service field.
Range: 0 through 255

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request.
Range: 0 through 128

vpn-interface *vpn-interface*—(Optional) Name of VPN interface over which to send packets.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

Sample Output: traceroute

```
user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2  red14  (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
 3  yellow (10.156.169.254)  7.737 ms  89.905 ms  0.834 ms
```

Sample Output: traceroute nnoresolve

```
user@host> traceoute santacruz noresolve
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

Sample Output: traceroute through an MPLS LSP

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
     MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

Sample Output: traceroute as-number-lookup (host)

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1  10.39.1.1 (10.39.1.1)  0.779 ms  0.728 ms  0.562 ms
 2  10.39.1.6 (10.39.1.6) [AS  32]  0.657 ms  0.611 ms  0.617 ms
 3  10.100.1.1 (10.100.1.1) [AS  10, 40, 50]  0.880 ms  0.808 ms  0.774 ms
```

Chapter 4
SNMP Monitoring and Troubleshooting

Table 8 summarizes the command-line interface (CLI) commands that allow you to monitor SNMP. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 8: Commands for Monitoring and Managing SNMP

Task Category	Task or Information to Monitor	Command
SNMP statistics	Display information about rmon alarms and events.	show snmp rmon on page 119
	Display statistics about SNMP packets sent and received	show snmp statistics on page 123
	Clear SNMP statistics.	clear snmp statistics on page 118
RMON	RMON alarm information.	show snmp rmon on page 119

clear snmp statistics

Syntax clear snmp statistics

Description Clear SNMP statistics.

Required Privilege Level clear

Sample Output

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290

user@host> clear snmp statistics

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 0, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0
```

show snmp rmon

Syntax show snmp rmon <alarms | events><brief | detail>>

Description Display information about rmon alarms and events.

Options none—Display information about all rmon alarms and events.

brief—(Optional) Display brief information about rmon alarms or events.

detail—(Optional) Display very detailed information about rmon alarms or events.

alarms—(Optional) Display information about rmon alarms.

events—(Optional) Display information about rmon events.

Required Privilege Level view

Options at a Glance Table 9 summarizes which information is included in the output of each show snmp rmon command options for each alarm and event. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 9: Show SNMP Rmon Output Field Summary

Options	Field Description
All	Alarm Index—Alarm identifier.
All	Event Index—Event identifier.
All	State—State of the alarm and/or event entry.
Alarms Option	
Detail	Creator— Mechanism by which the entry was configured.
Detail	Current value—Current value of the monitored variable in the most recent sample interval.
Detail	Falling event index—Identifies the event that is triggered when the falling threshold is crossed.
Detail	Falling threshold—Lower limit threshold value that is configured by the user.
Detail	Owner—Name of the entry that is configured by the user.
Detail	Rising event index—Identifies the event that is triggered when the rising threshold is crossed.
Detail	Rising threshold—Upper limit threshold value that is configured by the user.
Detail	Sample interval—Time period between samples is expressed in seconds.
Detail	Sample type—Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds.
Detail	Startup alarm—Alarm that may be sent when this entry is first activated.
Detail	Variable OID—Object ID to which the variable name is resolved.
All	Variable name—Particular SNMP object instance that is being monitored.
Events Option	
Detail	Creator— Mechanism by which the entry was configured.
Detail	Community—Identifies the trap group used for sending the SNMP trap.
Standard Brief	Last Event—Date and time of the last event's occurrence.
Standard Detail	Type—Type of notification made when an event is triggered.

Output Fields Alarm Index—Alarm identifier.

State—State of the alarm or event entry. It can be one of the following:

- Alarms
  - active—Entry is fully configured and activated.
  - falling threshold crossed—Value of the variable has crossed the lower threshold limit.
  - rising threshold crossed—Value of the variable has crossed the upper threshold limit.
  - under creation—Entry is being configured and is not yet activated.
  - startup—Alarm is waiting for the first sample of the monitored variable.
  - object not available—Monitored variable of that type is not available to the SNMP agent.
  - instance not available—Monitored variable's instance is not available to the SNMP agent.
  - object type invalid—Monitored variable is not a numeric value.
  - object processing errored—An error was encountered while processing the monitored variable.
  - unknown—State is not one of the above.
- Events
  - active—Entry has been fully configured and activated.
  - under creation—Entry is being configured and is not yet activated.
  - unknown—State is not one of the above.

Variable name—Name of the SNMP object instance being monitored.

Event Index—Event identifier.

Type—Type of notification made when an event is triggered. It can be one of the following:

- log—A system log message is generated and an entry is made to the log table.
- snmptrap—An SNMP trap is sent to the configured destination.
- log and trap—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination.
- none—Neither log or trap will be sent.

Last Event—Date and time of the last event. It has the format yyyy-mm-dd hh:mm:ss timezone.

Community—Identifies the trap group used for sending the SNMP trap.

Variable OID— Object ID to which the variable name is resolved. The format is “number. number.number.number. etcetera”.

Sample type—Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value.

Startup alarm—Alarm that might be sent when this entry is first activated, depending on the following criteria:

- Alarm is sent when
  - Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm.
  - Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm.
- Alarm is not sent when
  - Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm.
  - Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm.
  - Value of the alarm is between the thresholds.

Owner—Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.

Creator—Mechanism by which the entry was configured (CLI or SNMP).

Sample interval—Time period between samples (in seconds).

Rising threshold—Upper limit threshold value configured by the user.

Falling threshold—Lower limit threshold value configured by the user.

Rising event index—Identifies the event triggered when the rising threshold is crossed.

Falling event index—Identifies the event triggered when the falling threshold is crossed.

Current value—Current value of the monitored variable in the most recent sample interval.

Sample Output: show snmp rmon

```
user@host> show snmp rmon

Alarm
Index  State                        Variable name
    1  falling threshold crossed   ifInOctets.1

Event
Index  Type                         Last Event
    1  log and trap                2002-01-30 01:13:01 PST
```

Sample Output: show snmp rmon alarms

```
user@host> show snmp rmon alarms

Alarm
Index  State                        Variable name
    1  falling threshold crossed   ifInOctets.1
```

Sample Output: show snmp rmon alarms brief

```
user@host> show snmp rmon alarms brief

Alarm
Index  State                        Variable name
    1  falling threshold crossed   ifInOctets.1
```

Sample Output: show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail

Alarm Index 1:
  Variable name                        ifInOctets.1
  Variable OID                         1.3.6.1.2.1.2.2.1.10.1
  Sample type                          delta value
  Startup alarm                        rising or falling alarm
  Owner                                monitor
  Creator                              CLI
  State                                falling threshold crossed
  Sample interval                   60 seconds
  Rising threshold              100000
  Falling threshold              80000
  Rising event index                 1
  Falling event index                1
  Current value                      0
```

Sample Output: show snmp rmon events

```
user@host> show snmp rmon events

Event
Index  Type                         Last Event
    1  log and trap                2002-01-30 01:13:01 PST
```

Sample Output: show snmp rmon events brief

```
user@host> show snmp rmon events brief

Event
Index  Type                         Last Event
    1  log and trap                2002-01-30 01:13:01 PST
```

Sample Output: show snmp rmon events detail

```
user@host> show snmp rmon events detail

Event Index 1:
  Type                                 log and trap
  Community                            boy-elroy
  Last event                           2002-01-30 01:13:01 PST
  Creator                              CLI
  State                                active
```

show snmp statistics

Syntax show snmp statistics

Description Display statistics about SNMP packets sent and received by the router.

Output Fields Input—Information about received packets:

Packets—(snmpInPkts) Total number of messages delivered to the SNMP entity from the transport service.

Bad versions—(snmpInBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version.

Bad community names—(snmpInBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity.

Bad community uses—(snmpInBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message.

ASN parse errors—(snmpInASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.

Too bigs—(snmpInTooBigs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig.

No such names—(snmpInNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName.

Bad values—(snmpInBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue.

Read onlys—(snmpInReadOnlys) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr.

Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs.

Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs.

Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity.

Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity.

Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity.

Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity.

Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.

Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in such a way (other than a timeout) such that no response PDU could be returned.

Output—Information about transmitted packets:

Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service.

Too bigs—(snmpOutTooBigs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig.

No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName.

Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue.

General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated the SNMP entity with an error status field of genErr.

Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs that have been generated by the SNMP entity.

Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs that have been generated by the SNMP entity.

Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs that have been generated by the SNMP entity.

Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs that have been generated by the SNMP entity.

Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Sample Output

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290
```

Chapter 5

M-Series Chassis Monitoring and Troubleshooting

From the command-line interface (CLI), you can display information about the M-series chassis components using the commands in the show chassis hierarchy. These commands get their information from the chassis software process, which runs when the router is up and running. If the SCB, SSB, SFM, or FEB is not running, no information about chassis components is available through the CLI.

Table 10 summarizes the CLI commands you can use to monitor the M-series router chassis. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 10: Commands for Monitoring the M-series Chassis

Task or Information to Monitor	CLI Command
Chassis alarm status	show chassis alarms on page 135
Chassis clock-source configuration	request chassis pcg slot on page 128
Information currently on craft display	show chassis craft-interface on page 136
Environmental information	show chassis environment on page 138 show chassis feb on page 149 show chassis environment fpc on page 141 show chassis environment fpm on page 142 show chassis environment mcs on page 143 show chassis environment pcg on page 144 show chassis environment pem on page 145 show chassis environment routing-engine on page 146 show chassis environment sfm on page 147
Firmware version	show chassis firmware on page 150
FPC and PIC status	show chassis fpc on page 151 request chassis fpc slot on page 128 request chassis pic on page 129
Hardware inventory	show chassis hardware on page 153
MAC address	show chassis mac-addresses on page 157
Routing Engine	show chassis routing-engine on page 158 request chassis routing-engine master on page 130

Task or Information to Monitor	CLI Command
FEB, SCB, SFM, or SSB status	show chassis feb on page 149
	show chassis scb on page 162
	show chassis sfm on page 163 request chassis sfm slot on page 131 request chassis sfm master switch on page 132
	show chassis ssb on page 165 request chassis ssb master switch on page 133
Display a message on the router's craft interface	set chassis display message on page 134

request chassis fpc slot

Syntax request chassis fpc slot *slot number* <offline | online | restart>

Description (M40e and M160 routers only) Control the operation of the FPC.

Options offline—(Optional) Take the FPC offline.

online—(Optional) Put the FPC online.

restart—(Optional) Restart the FPC.

slot *slot number*—FPC slot number. It can be 0, 1, 2, 3, 4, 5, 6, or 7.

Required Privilege Level maintenance

request chassis pcg slot

Syntax request chassis pcg slot *slot number* <offline | online>

Description (M40e and M160 routers only) Control the operation of the PCG.

Options offline—Take the PCG offline.

online—(Optional) Put the PCG online.

slot *slot number*—(Optional) PCG slot number. It can be 0 or 1.

Required Privilege Level maintenance

request chassis pic

Syntax request chassis pic *fpc-slot pic-slot* <offline | online >

Description (M5, M10, M20 and M40 routers only) Control the operation of the PIC.

Options *fpc-slot*—FPC slot number. On M5 and M10 routers, it can be 0 or 1. On M20 routers, it can be 0, 1, 2, or 3. On M40 routers, it can be 0, 1, 2, 3, 4, 5, 6, or 7.

offline—(Optional) Take the PIC offline.

online—(Optional) Put the PIC online.

pic-slot—PIC slot number. It can be 0, 1, 2, or 3.

Required Privilege Level maintenance

request chassis routing-engine master

Syntax request chassis routing-engine master <acquire <no-confirm>| release | switch>

Description For routers with multiple Routing Engines only, control which Routing Engine is the master.

> **Warning**
>
> For routers that have two Routing Engines, both Routing Engines must be running JUNOS Release 4.0 or later. Do not run JUNOS Release 3.4 on one of the Routing Engines and Release 4.0 on the other. (Note that JUNOS Release 3.4 does not support Routing Engine redundancy, so if you are using this release of the software, only one Routing Engine can be installed in the router. It can be installed in either slot.)

By default, the Routing Engine in slot 0 (RE0) is the master and the Routing Engine in slot 1 (RE1) is the backup. To change the default master Routing Engine, include the routing-engine statement at the [edit chassis redundancy] hierarchy level in the configuration. For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started.*

To force the backup Routing Engine to become the master Routing Engine, use the request chassis routing-engine master command. If you use this command to change the master, and then restart the chassis software for any reason, the master reverts to the default setting.

> **Note**
>
> The configurations on the two Routing Engines do not have to be the same, and they are not automatically synchronized. If you configure both Routing Engines as masters, when the chassis software restarts for any reason, the Routing Engine in slot 0 becomes the master and the one in slot 1 becomes the backup.
>
> We recommend making both configurations the same.

When both Routing Engines have the same configuration and come up at the same time, the Routing Engine in slot 0 takes precedence over the one in slot 1. Table 11 indicates which Routing Engine takes the mastership based on the configuration of both Routing Engines.

Table 11: Routing Engine Mastership Election

		Slot 0 Configuration		
		Master	**Backup**	**Disabled**
Slot 1 Configuration	**Master**	Slot 0: master Slot 1: backup	Slot 0: backup Slot 1: master	Slot 0: disabled Slot 1: master
	Backup	Slot 0: master Slot 1: backup	Slot 0: master Slot 1: backup	Slot 0: disabled Slot 1: master
	Disabled	Slot 0: master Slot 1: disabled	Slot 0: master Slot 1: disabled	Slot 0: disabled Slot 1: disabled

Options acquire—(Optional) Attempt to become the master Routing Engine.

no-confirm—(Optional) Do not request confirmation.

release—(Optional) Request the other Routing Engine to become the master.

switch—(Optional) Toggle mastership between Routing Engines.

Required Privilege Level maintenance

See Also The routing-engine statement at the [edit chassis redundancy] hierarchy level in the configuration in the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service.*

Sample Output

```
user@M20-host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted
Reset other routing engine and become master ? [yes,no] (no)
```

Switch the master from M20-host-0 to the other Routing Engine:

```
root@m20-host-0> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The other Routing Engine becomes the master.
root@m20-host-0>
```

Have the backup Routing Engine M20-host-0 become the master:

```
root@m20-host-0> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.
root@m20-host-0>
```

request chassis sfm slot

Syntax request chassis sfm slot *slot number* < offline | online | restart >

Description (M40e and M160 routers only) Control the operation of an SFM.

Options offline—(Optional) Take the SFM offline.

online—(Optional) Put the SFM online.

restart—(Optional) Restart the SFM.

slot *slot number*—SFM slot number. It can be 0, 1, 2, or 3.

Required Privilege Level maintenance

request chassis sfm master switch

Syntax request chassis sfm master switch <no-confirm>

Description (M40e routers only) Control which SFM is master.

By default, the SFM in slot 0 (SFM0) is the master and the SFM in slot 1 (SFM1) is the backup.

To force the backup SFM to become the master SFM, use the request chassis sfm master switch command. If you use this command to change the master, and then restart the chassis software for any reason, the master reverts to the default setting.

To change the default master SFM, include the sfm statement at the [edit chassis redundancy] hierarchy level in the configuration. For more information, see the *JUNOS Internet Software Configuration Guide: Getting Started*.

Options no-confirm—Do not display a switch warning or query.

Required Privilege Level maintenance

Sample Output

```
user@host> request chassis sfm master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system forwarding module? [yes,no] (no) yes

Switch initiated, use “show chassis sfm” to verify
```

Sample Output: no-confirm option

```
user@host>request chassis sfm master switch no-confirm
Switch initiated, use “show chassis sfm” to verify
```

request chassis ssb master switch

Syntax request chassis ssb master switch <no-confirm>

Description (M20 routers only) Control which system switch board is master.

> **Warning** For routers that have two SSBs, both SSBs must be running JUNOS Release 4.0 or later. Do not run JUNOS Release 3.4 on one of the SSBs and Release 4.0 or later on the other. (Note that JUNOS Release 3.4 does not support SSB redundancy, so if you are using this release of the software, only one SSB can be installed in the router. It can be installed in either slot.)

By default, the SSB in slot 0 (SSB0) is the master and the SSB in slot 1 (SSB1) is the backup. To change the default master SSB, include the ssb statement at the [edit chassis redundancy] hierarchy level in the configuration. For more information, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service.*

To force the backup SSB to become the master SSB, use the request chassis ssb master switch command. If you use this command to change the master, and then restart the chassis software for any reason, the master reverts to the default setting.

> **Note** The configurations on the two SSBs do not have to be the same, and they are not automatically synchronized. If you configure both SSBs as masters, when the chassis software restarts for any reason, the SSB in slot 0 becomes the master and the one in slot 1 becomes the backup.

Options no-confirm—Will not display a warning and a query if it really should switch.

Required Privilege Level maintenance

Sample Output: request chassis ssb master switch

```
user@m20-host> request chassis ssb master switch
warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between system switch boards ? [yes,no] (no) yes

Switch initiated, use “show chassis ssb” to verify
```

Sample Output: request chassis ssb master switch (with no-confirm option)

```
user@m20-host>request chassis ssb master switch no-confirm
Switch initiated, use “show chassis ssb” to verify
```

set chassis display message

Syntax set chassis display message "*message*"

Description (M40, M40e, and M160 routers only) Display a text message on the craft interface display, which is on the front of the router. The craft interface alternates the display of text message and the standard craft interface messages, switching between messages every 2 seconds.

The text message is displayed for 5 minutes.

The craft interface display has four 20-character lines.

Options "*message*"—Message to display on the craft interface display. If the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. To stop a message, enter the command, omitting the *message* string.

Required Privilege Level clear

See Also show chassis craft-interface on page 136

Sample Output

```
user@sheep> set chassis display message "NOC contact Dusty (888) 526-1234"
message sent
user@sheep> show chassis craft-interface

Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off

FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green    .  .  *  .  .  *  *  .
Red      .  .  .  .  .  .  .  .

LCD screen:
     +--------------------+
     |NOC contact Dusty   |
     |(888) 526-1234      |
     |                    |
     |                    |
     +--------------------+


user@sheep> set chassis display message ""
message sent
user@sheep> show chassis craft-interface

Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
```

```
FPCs     0  1  2  3  4  5  6  7
------------------------------
Green    .  .  *  .  .  *  *  .
Red      .  .  .  .  .  .  .  .

LCD screen:
     +--------------------+
     |sheep               |
     |Up: 0+17:05:47      |
     |                    |
     |Temperature OK      |
     +--------------------+
```

show chassis alarms

Syntax show chassis alarms

Description Display information about the conditions that have been configured to trigger alarms.

It is not possible to clear the alarms for chassis components. This design is intentional; you must remedy the cause of the alarm. The alarm lights on the chassis are bright and are meant to be annoying. When they are lit, it indicates that you are running the router in a manner that we do not recommend.

You can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) nor extinguish the alarm LEDs. In addition, new alarms that occur after silencing an external device reactivate the external device.

Required Privilege Level view

Output Fields Alarm time—Date and time alarm was first recorded.

Class—Severity class for this alarm. It can be Minor or Major.

Description—Information about the alarm.

Sample Output: show chassis alarms (alarms active)

```
user@host> show chassis alarms
3 alarms are currently active
Alarm time Class Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed
```

Sample Output: show chassis alarms (no alarms active)

```
user@host> show chassis alarms
No alarms are currently active
```

show chassis craft-interface

Syntax show chassis craft-interface

Description For routers that have a display on the craft interface, show the messages that are currently displayed.

Required Privilege Level view

See Also set chassis display message on page 134

Output Fields FPM Display contents—Display contents of the Front Panel Module display.

router-name—The first line shows the name of the router.

Up—How long the router has been operational, in days, hours, minutes, and seconds.

message—Information about the router traffic load, the power supply status, the fan status, and the temperature status. The display of this information changes every 2 seconds.

If a text message has been created with the set chassis display command, this message appears on all four lines of the craft interface display. The display alternates between the text message and the standard system status messages every 2 seconds.

Front Panel System LEDs—Displays status of the Front Panel System LEDs. A dot (.) indicates the LED is not lit. An asterisk (*) indicates the LED is lit.

Front Panel Alarm Indicators—Displays status of the Front Panel Alarm Indicators. A dot (.) indicates the relay is off. An asterisk (*) indicates the relay is active.

Front Panel FPC LEDs—Displays status of the Front Panel FPC LEDs. A dot (.) indicates the LED is not lit. An asterisk (*) indicates the LED is lit.

MCS and SFM LEDs—Displays status of the MCS and SFM LEDs. A dot (.) indicates the LED is not lit. An asterisk (*) indicates the LED is lit. When neither a dot nor an asterisk is displayed, there is no board in that slot.

Sample Output: show chassis craft-interface (for M20 router)

```
user@sheep> show chassis craft-interface

Red alarm:     LED off, relay off
Yellow alarm:  LED on, relay on
Host OK LED:   On
Host fail LED: Off

FPCs     0  1  2  3
-------------------
Green    .  *  *  .
Red      .  .  .  .

LCD screen:\LCD Screen:
     +--------------------+
     |sheep               |
     |1 Alarm active      |
     |Y: FERF             |
     |                    |
     +--------------------+
```

Sample Output: show chassis craft-interface (for M40 router)

```
user@sheep> show chassis craft-interface
Front Panel LCD Display: enabled

Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host Fail LED: Off

NICs    0  1  2  3  4  5  6  7
-------------------------------
Green   *  .  *  .  *  .  *  .
Red     .  .  .  .  .  .  .  .

LCD Screen:
     +--------------------+
     |sheep               |
     |Up: 27+18:52:37     |
     |                    |
     |52.649kpps Load     |
     +--------------------+
```

Sample Output: show chassis craft-interface (for M160 router)

```
user@sheeps> show chassis craft-interface
FPM Display contents:
    +--------------------+
    |sheeps              |
    |Up: 1+16:46         |
    |                    |
    |Fans OK             |
    +--------------------+

Front Panel System LEDs:
Host     0    1
------------------------
OK       .    *
Fail     .    .
Master   .    *

Front Panel Alarm Indicators:
-----------------------------
Red LED      .
Yellow LED   .
Major relay  .
Minor relay  .

Front Panel FPC LEDs:
FPC    0   1   2   3   4   5   6   7
-----------------------------------
Red    .   .   .   .   .   .   .   .
Green  *   *   .   .   .   .   .   .

MCS and SFM LEDs:
MCS    0   1      SFM    0   1   2   3
--------------------------------------
Amber      .             .   .
Green      .             .   .
Blue       *             *   *
```

show chassis environment

Syntax show chassis environment

Description Display environmental information about the router chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

Required Privilege Level view

Output Fields Power—(M20 and M40 routers only) Information about each power supply. Status can be OK, Testing (during initial power-on), Failed, or Absent.
(M160 routers only) Information about the Power Entry Modules. Status can be OK, Testing (during initial power-on), Failed, or Absent.

Temp—Temperature of air flowing through the chassis.

Fans—Information about the fans. Status can be OK, Testing (during initial power-on), Failed, or Absent. Measurement indicates if spinning at normal or high speed.

Misc—Information about other components of the chassis. It could indicate the presence of one or more components.

On the M160 router, Misc includes CIP (Connector Interface Panel). OK indicates the CIP is present.

Sample Output: show chassis environment (for an M5 router)

```
user@m5-host> show chassis environment
Class Item                    Status     Measurement

Power Power Supply A          OK
      Power Supply B          OK

Temp  FPC Slot 0              OK         32 degrees C / 89 degrees F
      FEB                     OK         31 degrees C / 87 degrees F
      PS Intake               OK         26 degrees C / 78 degrees F
      PS Exhaust              OK         31 degrees C / 87 degrees F

Fans  Left Fan 1              OK         Spinning at normal speed
      Left Fan 2              OK         Spinning at normal speed
      Left Fan 3              OK         Spinning at normal speed
      Left Fan 4              OK         Spinning at normal speed

Misc  Craft Interface         OK
```

Sample Output: show chassis environment (for an M10 router)

```
user@m10-host> show chassis environment
Class Item                    Status     Measurement

Power Power Supply A          OK
      Power Supply B          OK

Temp  FPC Slot 0              OK         32 degrees C / 89 degrees F
      FPC Slot 1              OK         32 degrees C / 89 degrees F
      FEB                     OK         31 degrees C / 87 degrees F
      PS Intake               OK         26 degrees C / 78 degrees F
      PS Exhaust              OK         31 degrees C / 87 degrees F

Fans  Left Fan 1              OK         Spinning at normal speed
      Left Fan 2              OK         Spinning at normal speed
      Left Fan 3              OK         Spinning at normal speed
      Left Fan 4              OK         Spinning at normal speed

Misc  Craft Interface         OK
```

Sample Output: show chassis environment (for an M20 router)

```
user@m20-host> show chassis environment
Class Item                    Status     Measurement

Power Power Supply A          Absent
      Power Supply B          OK

Temp  FPC Slot 1              OK         37 degrees C / 98 degrees F
      FPC Slot 2              OK         38 degrees C / 100 degrees F
      Power Supply A          Absent
      Power Supply B          OK         37 degrees C / 98 degrees F
      SSB Slot 0              OK         36 degrees C / 96 degrees F
      Backplane               OK         32 degrees C / 89 degrees F

Fans  Rear Fan                OK         Spinning at normal speed
      Upper Fan               OK         Spinning at normal speed
      Middle Fan              OK         Spinning at normal speed
      Bottom Fan              OK         Spinning at normal speed

Misc  Craft Interface         OK
```

Sample Output: show chassis environment (for an M40 router)

```
user@m40-host> show chassis environment
Class Item                    Status     Measurement

Power Power Supply A          OK
      Power Supply B          Absent

Temp  FPC Slot 0              OK         34 degrees C / 93 degrees F
      FPC Slot 2              OK         32 degrees C / 89 degrees F
      FPC Slot 4              OK         31 degrees C / 87 degrees F
      FPC Slot 5              OK         39 degrees C / 102 degrees F
      SCB                     OK         28 degrees C / 82 degrees F
      Backplane @ A1          OK         35 degrees C / 95 degrees F
      Backplane @ A2          OK         32 degrees C / 89 degrees F

Fans  Top Impeller            OK         Spinning at normal speed
      Bottom Impeller         OK         Spinning at normal speed
      Rear Fan 1              OK         Spinning at normal speed
      Rear Fan 2              OK         Spinning at normal speed
      Rear Fan 3              OK         Spinning at normal speed

Misc  Craft Interface         OK
```

Sample Output: show chassis environment (for an M40e router)

```
user@m40e-host> show chassis environment
Class Item                    Status     Measurement
Power PEM 0                   OK
      PEM 1                   Check
Temp  PCG 0                   OK         40 degrees C / 104 degrees F
      PCG 1                   OK         42 degrees C / 107 degrees F
      Host 0                  OK         31 degrees C / 87 degrees F
      Host 1                  Absent
      MCS 0                   OK         44 degrees C / 111 degrees F
      MCS 1                   Absent
      SPP 0                   OK         41 degrees C / 105 degrees F
      SPR 0                   OK         45 degrees C / 113 degrees F
      SPP 1                   OK         44 degrees C / 111 degrees F
      SPR 1                   OK         45 degrees C / 113 degrees F
      FPC 0                   OK         42 degrees C / 107 degrees F
      FPC 1                   OK         46 degrees C / 114 degrees F
      FPC 2                   OK         41 degrees C / 105 degrees F
      FPC 4                   OK         41 degrees C / 105 degrees F
      FPC 6                   OK         38 degrees C / 100 degrees F
      FPM CMB                 OK         28 degrees C / 82 degrees F
      FPM Display             OK         28 degrees C / 82 degrees F
Fans  Rear Bottom Blower      OK         Spinning at normal speed
      Rear Top Blower         OK         Spinning at normal speed
      Front Top Blower        OK         Spinning at normal speed
      Fan Tray Rear Left      OK         Spinning at normal speed
      Fan Tray Rear Right     OK         Spinning at normal speed
      Fan Tray Front Left     OK         Spinning at normal speed
      Fan Tray Front Right    OK         Spinning at normal speed
Misc  CIP
```

Sample Output: show chassis environment (for an M160 router)

```
user@m160-host> show chassis environment
Class Item                    Status     Measurement
Power PEM 0                   OK
      PEM 1                   OK
Temp  PCG 0                   OK         41 degrees C / 105 degrees F
      PCG 1                   OK         49 degrees C / 120 degrees F
      Host 0                  Absent
      Host 1                  OK         37 degrees C / 98 degrees F
      MCS 0                   Absent
      MCS 1                   OK         56 degrees C / 132 degrees F
      SPP 0                   OK         40 degrees C / 104 degrees F
      SPR 0                   OK         44 degrees C / 111 degrees F
      SPP 1                   OK         43 degrees C / 109 degrees F
      SPR 1                   OK         47 degrees C / 116 degrees F
      FPC 0                   OK         41 degrees C / 105 degrees F
      FPC 1                   OK         44 degrees C / 111 degrees F
      FPM CMB                 OK         33 degrees C / 91 degrees F
      FPM Display             OK         34 degrees C / 93 degrees F
Fans  Rear Bottom Blower      OK         Spinning at high speed
      Rear Top Blower         OK         Spinning at high speed
      Front Top Blower        OK         Spinning at high speed
      Fan Tray Front Left     OK         Spinning at high speed
      Fan Tray Front Right    OK         Spinning at high speed
      Fan Tray Rear Left      OK         Spinning at high speed
      Fan Tray Rear Right     OK         Spinning at high speed
Misc  CIP                     OK
```

show chassis environment fpc

Syntax show chassis environment fpc <*slot*>

Description (M40e and M160 routers only) Display environmental information about the FPC in the router.

Options none—Display environmental information about all FPCs.

slot—(Optional) Display environmental information about an individual FPC. *slot* can be a value from 0 through 7.

Required Privilege Level view

Output Fields State—Status of the FPC. It can be Unknown, Empty, Present, Ready, Announce online, Online, Offline, or Diagnostics.

Temperature—Temperature of the air flowing past the FPC.

Power—Information about power supplies on the FPC.

CMB Revision—Revision level of the chassis management bus slave.

Sample Output

```
user@m160-host> show chassis environment fpc
FPC 0 status:
  State                       Online
  Temperature               42 degrees C / 107 degrees F
  Power:
    1.5 V                 1500 mV
    2.5 V                 2509 mV
    3.3 V                 3308 mV
    5.0 V                 4991 mV
    5.0 V bias            4952 mV
    8.0 V bias            8307 mV
  CMB Revision              12
FPC 1 status:
  State                       Online
  Temperature               45 degrees C / 113 degrees F
  Power:
    1.5 V                 1498 mV
    2.5 V                 2501 mV
    3.3 V                 3319 mV
    5.0 V                 5020 mV
    5.0 V bias            5025 mV
    8.0 V bias            8307 mV
  CMB Revision              12
```

show chassis environment fpm

Syntax show chassis environment fpm

Description (M40e and M160 routers only) Display environmental information about the front panel module (craft interface) in the router.

Required Privilege Level view

Output Fields State—Status of the front panel module. It can be Online or Offline.

FPM CMB Voltage—Information about the power supplied to the FPM Chassis Management Bus (CMB) cord.

FPM Display Voltage—Information about the FPM display power supply.

FPM CMB temperature—Temperature of the air flowing past the FPM CMB.

FPM Display temperature—Temperature of the air flowing past the FPM display.

CMB Revision—Revision level of the chassis management bus slave.

Sample Output

```
user@m160-host> show chassis environment fpm
FPM status:
  State                       Online
  FPM CMB Voltage:
    5.0 V bias           5030 mV
    8.0 V bias           8083 mV
  FPM Display Voltage:
    5.0 V bias           4998 mV
  FPM CMB temperature      34 degrees C / 93 degrees F
  FPM Display temperature  35 degrees C / 95 degrees F
  CMB Revision             12
```

show chassis environment mcs

Syntax show chassis environment mcs <*slot*>

Description (M40e and M160 routers only) Display environmental information about the Miscellaneous Control Subsystem (MCS).

Options none—Display environmental information about both MCSs.

slot—(Optional) Display environmental information about an individual MCS. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields State—Status of the MCS. It can be Present, Online, Offline, or Empty. Also indicates Master.

Temperature—Temperature of the air flowing past the MCS.

Power—Information about MCS power supplies.

FPGA Revision—Field Programmable Gate Array revision information.

CMB Revision—Revision level of the chassis management bus slave.

Sample Output

```
user@m160-host> show chassis environment mcs
MCS 0 status:
  State                           Online Master
  Temperature                   0 degrees C / 32 degrees F
  Power:
    3.3 V                    3318 mV
    5.0 V                    5001 mV
    12.0 V                  11833 mV
    5.0 V bias               4991 mV
    8.0 V bias               8341 mV
  CMB Revision                 12
  FPGA Revision                12
MCS 1 status:
  State                           Present
  Power:
    3.3 V                    3308 mV
    5.0 V                    5013 mV
    12.0 V                  11809 mV
    5.0 V bias               4952 mV
    8.0 V bias               8346 mV
  CMB Revision                 12
```

show chassis environment pcg

Syntax show chassis environment pcg <*slot*>

Description (M40e and M160 routers only) Display environmental information about the PFE Clock Generator (PCG).

Options none—Display environmental information about both PCGs.

slot—(Optional) Display environmental information about an individual PCG. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields *PCG slot*—Slot number. It can be 0 or 1.

State—Status of PCG. It can be Present, Online, Offline, or Empty. If Online, it can be the current PFE clock source or backup.

Temperature—Temperature of the air flowing past the PCG.

Frequency—Frequency setting and measurement for the PCG.

Power—Information about PCG power supplies.

CMB Revision—Revision level of the chassis management bus slave.

Sample Output

```
user@m160-host> show chassis environment pcg
PCG 0 status:
  State                       Online - PFE clock source
  Temperature              42 degrees C / 107 degrees F
  Frequency:
    Setting               125 MHz
    Measurement           125 MHz
  Power:
    3.3 V                3267 mV
    5.0 V bias           4932 mV
    8.0 V bias           8224 mV
  CMB Revision             12
PCG 1 status:
  State                       Online
  Temperature              49 degrees C / 120 degrees F
  Frequency:
    Setting               125 MHz
    Measurement           125 MHz
  Power:
    3.3 V                3264 mV
    5.0 V bias           4967 mV
    8.0 V bias           8236 mV
  CMB Revision             12
```

show chassis environment pem

Syntax show chassis environment pem <*slot*>

Description (M40e and M160 routers only) Display environmental information about the Power Entry Module (PEM).

Options none—Display environmental information about both PEMs.

slot—(Optional) Display environmental information about an individual PEM. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields slot—Name of PEM slot.

State—Status of PEM.

Temperature—Temperature of the air flowing past the PEM.

DC Input—Status of the DC input.

DC Output—Status of the DC output.

Load—Information about the load on supply, in percentage of rated current being used.

Voltage—Information about voltage of the PEM.

Sample Output

```
user@m160-host> show chassis environment pcg
PEM 0 status:
  State                       Online
  Temperature                 OK
  DC input                    OK
  DC output                   OK
  Load                        Less than 20 percent
  Voltage:
    48.0 V input        69028 mV
    48.0 V fan supply   48839 mV
    5.0 V bias           5013 mV
    8.0 V bias           8253 mV
PEM 1 status:
  State                       Online
  Temperature                 OK
  DC input                    OK
  DC output                   OK
  Load                        Less than 20 percent
  Voltage:
    48.0 V input        69307 mV
    48.0 V fan supply   49170 mV
    5.0 V bias           4991 mV
    8.0 V bias           8263 mV
```

show chassis environment routing-engine

Syntax show chassis environment routing-engine <*slot*>

Description (M40e and M160 routers only) Display environmental information about the Routing Engine.

Options none—Display environmental information about all Routing Engines.

slot—(Optional) Display environmental information about an individual Routing Engine. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields Routing engine slot—(For systems with multiple Routing Engines) Number of the Routing Engine slot. It can be 0 or 1.

State—Status of Routing Engine. It can be Present or Empty. If Present it can be Master.

Temperature—Temperature of the air flowing past the Routing Engine.

Sample Output

```
user@m160-host> show chassis environment routing-engine
Route engine 0 status:
  State:                      Present Master
  Temperature:              0 degrees C / 32 degrees F
Route engine 1 status:
  State:                      Present
```

show chassis environment sfm

Syntax show chassis environment sfm <*slot*>

Description (M40e and M160 routers only) Display environmental information about the Switching and Forwarding Module (SFM).

Options none—Display environmental information about all SFMs.

slot—(Optional) Display environmental information about an individual SFM. *slot* can be 0 through 3.

Required Privilege Level view

Output Fields *sfm plane*—Slot number of SFM. It can be 0, 1, 2, or 3.

State—Status of the SFM.

SPP Temperature—Temperature of the air flowing past the Switch Plane Processor card.

SPR Temperature—Temperature of the air flowing past the Switch Plane Router.

SPP Power—Information about Switch Plane Processor card power supplies.

SPR Power—Information about Switch Plane Router power supplies.

CMB Revision—Chassis Management Bus slave revision identification.

Sample Output

```
user@m160-host> show chassis environment sfm 0
SFM 0 status:
  State                       Online
  SPP temperature          40 degrees C / 104 degrees F
  SPR temperature          44 degrees C / 111 degrees F
  SPP Power:
    1.5 V                1504 mV
    2.5 V                2479 mV
    3.3 V                3285 mV
    5.0 V                5028 mV
    5.0 V bias           4967 mV
  SPR Power:
    1.5 V                1501 mV
    2.5 V                2485 mV
    3.3 V                3291 mV
    5.0 V                5023 mV
    5.0 V bias           4967 mV
    8.0 V bias           8351 mV
  CMB Revision             12
```

```
SFM 1 status:
  State                       Online
  SPP temperature          44 degrees C / 111 degrees F
  SPR temperature          47 degrees C / 116 degrees F
  SPP Power:
    1.5 V                1498 mV
    2.5 V                2494 mV
    3.3 V                3293 mV
    5.0 V                5020 mV
    5.0 V bias           5001 mV
  SPR Power:
    1.5 V                1500 mV
    2.5 V                2496 mV
    3.3 V                3296 mV
    5.0 V                5013 mV
    5.0 V bias           4998 mV
    8.0 V bias           8358 mV
  CMB Revision             12
```

show chassis feb

Syntax show chassis feb

Description (M5 and M10 routers only) Display status information about the FEB (Forwarding Engine Board).

Required Privilege Level view

Output Fields Temperature—Temperature of the air flowing past the FEB.

CPU utilization—Percentage of CPU being used by the FEB's processor.

Interrupt utilization—Of the total CPU being used by the FEB's processor, the percentage being used for interrupts.

Heap Utilization—Percentage of heap space (dynamic memory) being used by the FEB's processor. If this number exceeds 80%, it might indicate a software problem (memory leak).

Buffer Utilization—Percentage of buffer space being used by the FEB's processor for buffering internal messages.

DRAM—Amount of DRAM available to the FEB's CPU.

Internet Processor II—Information about the FEB processor.

Start time—Time when the Routing Engine noticed that the FEB was running.

Uptime—How long the Routing Engine has been connected to the FEB and, therefore, how long the FPC has been up and running.

Sample Output

```
user@host> show chassis feb
FEB status:
  Temperature                 37 Centigrade
  CPU utilization              0 percent
  Interrupt utilization        0 percent
  Heap utilization            16 percent
  Buffer utilization          43 percent
  DRAM                        64 Mbytes
  Internet Processor II          Version 1, Foundry IBM, Part number 9
  Start time                     1999-01-24 16:24:42 UTC
  Uptime                         2 hours, 21 minutes, 28 seconds
```

show chassis firmware

Syntax show chassis firmware

Description Display the version levels of the firmware running on the SCB, SFM, SSB, FEB, and FPCs.

Required Privilege Level view

Output Fields Part—Name of chassis part.

Type—Type of firmware. It can be ROM or O/S.

Version—Version of firmware runing on the chassis part.

Sample Output

```
user@m10-host> show chassis firmware
Part                     Type       Version
Forwarding engine board  ROM        Juniper ROM Monitor Version 4.1b2
                         O/S        Version 4.1I1 by tlim on 2000-04-24 11:27

user@m20-host> show chassis firmware
Part                     Type       Version
System switch board      ROM        Juniper ROM Monitor Version 3.4b26
                         O/S        Version 3.4I16 by smackie on 2000-02-29 2
FPC 1                    ROM        Juniper ROM Monitor Version 3.0b1
                         O/S        Version 3.4I4 by smackie on 2000-02-25 21
FPC 2                    ROM        Juniper ROM Monitor Version 3.0b1
                         O/S        Version 3.4I4 by smackie on 2000-02-25 21

user@m40-host> show chassis firmware
Part                     Type       Version
System control board     ROM        Juniper ROM Monitor Version 2.0i126Copyri
                         O/S        Version 2.0i1 by root on Thu Jul 23 00:51
FPC 5                    ROM        Juniper ROM Monitor Version 2.0i49Copyrig
                         O/S        Version 2.0i1 by root on Thu Jul 23 00:59
user@m160-host> show chassis firmware
Part                     Type       Version
SFM 0                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by tlim on 2000-02-29 11:50
SFM 1                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by tlim on 2000-02-29 11:50
FPC 0                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 1                    ROM        Juniper ROM Monitor Version 4.0b2
                         O/S        Version 4.0I1 by tlim on 2000-02-29 11:56
FPC 2                    ROM        Juniper ROM Monitor Version 4.0b3
                         O/S        Version 4.0I1 by tlim on 2000-02-29 11:56
```

show chassis fpc

Syntax show chassis fpc <pic-status <*fpc-slot*>> <detail <*fpc-slot*>>

Description Display status information about the installed FPCs and PICs.

Options none—Display brief information about the state of the FPCs.

detail <*fpc-slot*>—(Optional) Display detailed status information for the all FPCs or for the specified FPC.

pic-status <*fpc-slot*>—(Optional) Display information for all PICs or for the PICs in the specified slot.

Required Privilege Level view

Output Fields Slot—Slot number and state. The state can be one of the following:

- dead—Held in reset because of errors.
- diag—Slot is being ignored while the FPC is running diagnostics.
- dormant—Held in reset.
- empty—No FPC is present.
- online—FPC is online and running.
- probed—Probe is complete; awaiting restart of PFE.
- probe-wait—Waiting to be probed.

Logical slot—Slot number.

Temp (C) or Temperature—Temperature of the air passing by the FPC, in degrees Centigrade.

Total CPU Utilization (%)—(Standard output only) Total percentage of CPU being used by the FPC's processor.

Interrupt CPU Utilization (%)—(Standard output only) Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.

Memory DRAM—(Standard output only) Total DRAM available to the FPC's processor, in megabytes.

Heap Utilization (%)—(Standard output only) Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80%, it might indicate a software problem (memory leak).

Buffer Utilization (%)—(Standard output only) Percentage of buffer space being used by the FPC's processor for buffering internal messages.

Total CPU DRAM—(Detail output only) Amount of DRAM available to the FPC's CPU.

Total SRAM—(Detail output only) Amount of SRAM used by the FPC's CPU.

Total SDRAM—(Detail output only) Total amount of memory used for storing packets and notifications.

Total notification SDRAM—(Detail output only) Amount of memory used by the Packet Forwarding Engine for packet buffer and packet notification space.

I/O Manager ASIC information—(Detail output only) For the I/O Manager, identifies the version number, manufacturer, and part number.

Start time—(Detail output only) Time when the Routing Engine noticed that the FPC was running.

Uptime—(Detail output only) 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 PIC at each PIC location.

Sample Output: show chassis fpc

```
user@m10-host> show chassis fpc
FPC status:
                Temp
Slot State       (C)
  0  Online       27
  1  Online       28

user@m20-host> show chassis fpc
FPC status:
                Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State       (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Empty         0      0          0         0         0          0
  1  Online       38      0          0         8         0          4
  2  Online       35      0          0         8         0          3
  3  Empty         0      0          0         0         0          0
```

Sample Output: show chassis fpc detail

```
user@host> show chassis fpc detail 1
Slot 1 information:
  State                               Online
  Temperature                      48 degrees C
  Total CPU DRAM                   32 Mbytes
  Total SRAM                        4 Mbytes
  Total SDRAM                     256 Mbytes
  I/O Manager ASICs information       Version 2.0, Foundry IBM, Part number 0
  I/O Manager ASICs information       Version 2.0, Foundry IBM, Part number 0
  Start time                          2000-02-08 02:18:49 UTC
  Uptime                              14 hours, 41 minutes, 41 seconds
```

Sample Output: show chassis fpc pic-status

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1    1x OC-12 ATM, MM
  PIC 2    1x OC-12 ATM, MM
  PIC 3    1x OC-12 ATM, MM
Slot 1 Online
  PIC 0    1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0    1x OC-192 SONET, SMSR
```

show chassis hardware

Syntax show chassis hardware <detail>

Description Display a list of all FPCs and PICs installed in the router chassis, including the hardware version level and serial number.

Options none—Display information about hardware.

detail—(Optional) Display detailed information about hardware, including memory, the hardware version level, serial number, and additional information about memory.

Required Privilege Level view

Output Fields Item—Chassis component. Information is displayed about the backplane, the power supplies, the maxicab (the connection between the Routing Engine and the backplane), the SCB, SSB, SFM, or FEB and each of the FPCs and PICs.

Version—Revision level of the chassis component.

Part number—Part number of the chassis component.

Serial number—Serial number of the chassis component. The serial number of the backplane is also the serial number of the router chassis. Use this serial number when you need to contact technical support about the router chassis.

Description—Brief description of hardware item.

- For the power supplies, the type of supply.
- For the PICs, the type of PIC.
- For FPC, the type of FPC. It can be FPC Type 1, FPC Type 2, or FPC Type OC192.
- For host, the Routing Engine type.

Sample Output: show chassis hardware (standard, for an M10 router)

```
user@m10-host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                1122              M10
Midplane         REV 1.1  710-001950   S/N AC6626
Power supply A   Rev 01   740-002497   S/N LC36095       AC
Power supply B   Rev 01   740-002497   S/N LC36100       AC
Display          REV 1.2  710-001995   S/N AC6656
Host                                   18000005dfb3fb01  teknor
FEB              REV 01   710-001948   S/N AC6632        Internet Processor II
FPC 0
  PIC 0          REV 08   750-001072   S/N AB2485        1x G/E, 1000 BASE-SX
  PIC 1          REV 01   750-000613   S/N AA1048        1x OC-12 SONET, SMIR
FPC 1
```

Sample Output: show chassis hardware (for an M20 router)

```
user@m20-host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                20033             M20
Backplane        REV 07   710-001517   S/N AA7940
Power supply B   Rev 01   740-001465   S/N 000001        AC
Display          REV 02   710-001519   S/N AA9704
Host 0                                 98000004f8f27501  teknor
SSB slot 0       REV 01   710-001951   S/N AD5905        Internet Processor II
  SSRAM bank 0   REV 01   710-001385   S00480            2 Mbytes
  SSRAM bank 1   REV 01   710-001385   S00490            2 Mbytes
  SSRAM bank 2   REV 01   710-001385   S001:?            2 Mbytes
  SSRAM bank 3   REV 01   710-001385   S00483            2 Mbytes
SSB slot 1       N/A      N/A          N/A               Backup
FPC 1            REV 01   710-001292   S/N AB7528
  SSRAM          REV 01   710-000077   S/N 304209        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N 000603        64 Mbytes
  SDRAM bank 1   REV 01   710-000099   S/N 000414        64 Mbytes
  PIC 0          REV 03   750-000612   S/N AB8433        2x OC-3 ATM, MM
  PIC 1          REV 01   750-000616   S/N AA1168        1x OC-12 ATM, MM
  PIC 2          REV 01   750-000613   S/N AA1008        1x OC-12 SONET, SMIR
  PIC 3          REV 01   750-002501   S/N AD5810        4x E3
FPC 2            REV 01   710-001292   S/N AC0119
  SSRAM          REV 01   710-000077   S/N 503241        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N 306835        64 Mbytes
  SDRAM bank 1   REV 01   710-000099   S/N 306832        64 Mbytes
```

Sample Output: show chassis hardware (for an M40 router)

```
user@m40-host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Backplane        REV 02   710-000073   S/N AA0053
Power supply A   Rev 2    740-000235   S/N 000042        DC
Maxicab          REV X1   710-000229   S/N AAO139
Minicab          REV X1   710-000482   S/N AA0201
Display          REV 06   710-000150   S/N AA0905
Host                                                     cpv5000
SCB              REV X1   710-000075   S/N AA0158        Internet Processor I
  SSRAM bank 0   REV 02   710-000077   S/N AA2267        1 Mbyte
  SSRAM bank 1   REV 02   710-000077   S/N AA2270        1 Mbyte
  SSRAM bank 2   REV 02   710-000077   S/N AA2269        1 Mbyte
  SSRAM bank 3   REV 02   710-000077   S/N AA2268        1 Mbyte
FPC 0            REV 01   710-000175   S/N AA0048
  SSRAM          REV 01   710-000077   S/N AA2333        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N AA2332        64 Mbytes
  SDRAM bank 1   REV X1   710-000099   S/N AA2337        64 Mbytes
  PIC 0          REV 04   750-000613   S/N aa0343        1x OC-12 SONET, SMIR
  PIC 1          REV 04   750-000613   S/N AA0379        1x OC-12 SONET, SMIR
  PIC 2          REV 04   750-000613   S/N AA0377        1x OC-12 SONET, SMIR
  PIC 3          REV 04   750-000613   S/N AA0378        1x Tunnel
FPC 2            REV 01   710-000175   S/N AA0042
  SSRAM          REV 02   710-000077   S/N AA2288        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N AA2331        64 Mbytes
  SDRAM bank 1   REV 01   710-000099   S/N AA2330        64 Mbytes
  PIC 0          REV X1   750-000603   S/N AA0143        4x OC-3 SONET, SMIR
  PIC 1          REV X1   750-000615   S/N AA0149        4x OC-3 SONET, MM
  PIC 2          REV X1   750-000611   S/N AA0148        4x OC-3 SONET, MM
  PIC 3          REV 04   750-000613   S/N AA0330        1x OC-12 SONET, SMIR
```

```
FPC 4            REV 01   710-000175   S/N AA0050
  SSRAM          REV 01   710-000077   S/N AA2327        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N AA2329        64 Mbytes
  SDRAM bank 1   REV 01   710-000099   S/N AA2328        64 Mbytes
  PIC 0          REV 04   750-000613   S/N AA0320        1x OC-12 SONET, SMIR
  PIC 2          REV 05   750-000616   S/N AA1341        1x OC-12 ATM, MM
  PIC 3          REV 08   750-001072   S/N AB2462        1x G/E, 1000 BASE-SX
FPC 5            REV 10   710-000175   S/N AA7663
  SSRAM          REV 01   710-000077   S/N 501590        1 Mbyte
  SDRAM bank 0   REV 01   710-000099   S/N 300949        64 Mbytes
  SDRAM bank 1   REV 01   710-000099   S/N 300868        64 Mbytes
  PIC 1          REV 01   750-001323   S/N AB1670        1x Tunnel
```

Sample Output: show chassis hardware (for an M40e router)

```
user@m40e-host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                                  m40e
Midplane         REV 01   710-005071   AX3671
FPM CMB          REV 03   710-001642   AR9074
FPM Display      REV 03   710-001647   AR7331
CIP              REV 04   710-002649   BB4449
PEM 0            Rev 01   740-003787   MC12364           Power Entry Module
PEM 1            Rev 01   740-003787   MC12383           Power Entry Module
PCG 0            REV 07   710-001568   AG1332
PCG 1            REV 07   710-001568   AR3789
Host 0                                 3e000007c8176601  Present
MCS 0            REV 11   710-001226   AN5813
SFM 0 SPP        REV 07   710-001228   AG4676
SFM 0 SPR        REV 05   710-002189   AE4735            Internet Processor II
SFM 1 SPP        REV 07   710-001228   AP1347
SFM 1 SPR        REV 05   710-002189   BE0063            Internet Processor II
FPC 0            REV 02   710-005197   BD9845            M40e-FPC Type 2
  CPU            REV 01   710-004600   BD9504
  PIC 0          REV 03   750-003737   AY3991            4x G/E, 1000 BASE-SX
FPC 1            REV 01   710-005197   BD9842            M40e-FPC Type 2
  CPU            REV 01   710-004600   BB4869
  PIC 0          REV 07   750-001900   AR8278            1x OC-48 SONET, SMSR
FPC 2            REV 02   710-005197   BD9824            M40e-FPC Type 2
  CPU            REV 01   710-004600   BD9531
  PIC 0          REV 03   750-003737   AY3986            4x G/E, 1000 BASE-SX
FPC 4            REV 02   710-005078   BE0664            M40e-FPC Type 1
  CPU            REV 01   710-004600   BD9559
  PIC 0          REV 03   750-001894   AG7963            1x G/E, 1000 BASE-SX
  PIC 2          REV 01   750-002575   AF2472            4x OC-3 SONET, SMIR
FPC 6            REV 02   710-005078   BE0652            M40e-FPC Type 1
  CPU            REV 01   710-004600   BD9607
  PIC 0          REV 02   750-002911   AN2286            4x F/E, 100 BASE-TX
  PIC 2          REV 01   750-002577   AP6345            4x OC-3 SONET, MM
```

Sample Output: show chassis hardware (for an M160 router)

```
user@m160-host> show chassis hardware
Item             Version  Part number  Serial number     Description
Chassis                                101               M160
Midplane         REV 02   710-001245   S/N AB4107
FPM CMB          REV 01   710-001642   S/N AA2911
FPM Display      REV 01   710-001647   S/N AA2999
CIP              REV 02   710-001593   S/N AA9563
PEM 0            Rev 01   740-001243   S/N KJ35769       DC
PEM 1            Rev 01   740-001243   S/N KJ35765       DC
PCG 0            REV 01   710-001568   S/N AA9794
PCG 1            REV 01   710-001568   S/N AA9804
Host 1                                 da000004f8d57001  teknor
MCS 1            REV 03   710-001226   S/N AA9777
SFM 0 SPP        REV 04   710-001228   S/N AA2975
```

```
SFM 0 SPR        REV 02   710-001224   S/N AA9838        Internet Processor I
SFM 1 SPP        REV 04   710-001228   S/N AA2860
SFM 1 SPR        REV 01   710-001224   S/N AB0139        Internet Processor I
FPC 0            REV 03   710-001255   S/N AA9806        FPC Type 1
  CPU            REV 02   710-001217   S/N AA9590
  PIC 1          REV 05   750-000616   S/N AA1527        1x OC-12 ATM, MM
  PIC 2          REV 05   750-000616   S/N AA1535        1x OC-12 ATM, MM
  PIC 3          REV 01   750-000616   S/N AA1519        1x OC-12 ATM, MM
FPC 1            REV 02   710-001611   S/N AA9523        FPC Type 2
  CPU            REV 02   710-001217   S/N AA9571
  PIC 0          REV 03   750-001900   S/N AA9626        1x STM-16 SDH, SMIR
  PIC 1          REV 01   710-002381   S/N AD3633        2x G/E, 1000 BASE-SX
FPC 2                                                   FPC Type OC192
  CPU            REV 03   710-001217   S/N AB3329
  PIC 0          REV 01                                  1x OC-192 SM SR-2
```

Sample Output: show chassis hardware detail

```
user@host> show chassis hardware detail
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                101               M160
Midplane         REV 02   710-001245   S/N AB4107
FPM CMB          REV 01   710-001642   S/N AA2911
FPM Display      REV 01   710-001647   S/N AA2999
CIP              REV 02   710-001593   S/N AA9563
PEM 0            Rev 01   740-001243   S/N KJ35769       DC
PEM 1            Rev 01   740-001243   S/N KJ35765       DC
PCG 0            REV 01   710-001568   S/N AA9794
PCG 1            REV 01   710-001568   S/N AA9804
Host 1                                 da000004f8d57001  teknor
MCS 1            REV 03   710-001226   S/N AA9777
SFM 0 SPP        REV 04   710-001228   S/N AA2975
SFM 0 SPR        REV 02   710-001224   S/N AA9838        Internet Processor I
  SSRAM bank 0   REV 01   710-000077   S/N 306456        1 Mbyte
  SSRAM bank 1   REV 01   710-000077   S/N 306474        1 Mbyte
  SSRAM bank 2   REV 01   710-000077   S/N 306388        1 Mbyte
  SSRAM bank 3   REV 01   710-000077   S/N 306392        1 Mbyte
SFM 1 SPP        REV 04   710-001228   S/N AA2860
SFM 1 SPR        REV 01   710-001224   S/N AB0139        Internet Processor I
  SSRAM bank 0   REV 01   710-000077   S/N 302917        1 Mbyte
  SSRAM bank 1   REV 01   710-000077   S/N 302662        1 Mbyte
  SSRAM bank 2   REV 01   710-000077   S/N 302593        1 Mbyte
  SSRAM bank 3   REV 01   710-000077   S/N 100160        1 Mbyte
FPC 0            REV 03   710-001255   S/N AA9806        FPC Type 1
  CPU            REV 02   710-001217   S/N AA9590
  SSRAM          REV 01   710-000077   S/N 302836        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00141            32 Mbytes
  SDRAM 1        REV 01   710-001196   S0010;            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 302633        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00143            32 Mbytes
  SDRAM 1        REV 01   710-001196   S00115            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 302952        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00135            32 Mbytes
  SDRAM 1        REV 01   710-001196   S001=3            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 302892        1 Mbyte
  SDRAM 0        REV 01   710-001196   S000?6            32 Mbytes
  SDRAM 1        REV 01   710-001196   S001=5            32 Mbytes
  PIC 1          REV 05   750-000616   S/N AA1527        1x OC-12 ATM, MM
  PIC 2          REV 05   750-000616   S/N AA1535        1x OC-12 ATM, MM
  PIC 3          REV 01   750-000616   S/N AA1519        1x OC-12 ATM, MM
```

```
FPC 1            REV 02   710-001611   S/N AA9523        FPC Type 2
  CPU            REV 02   710-001217   S/N AA9571
  SSRAM          REV 01   710-000077   S/N 306340        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00012            32 Mbytes
  SDRAM 1        REV 01   710-001196   S0001?            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 306454        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00028            32 Mbytes
  SDRAM 1        REV 01   710-001196   S0002?            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 306492        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00015            32 Mbytes
  SDRAM 1        REV 01   710-001196   S00031            32 Mbytes
  SSRAM          REV 01   710-000077   S/N 306363        1 Mbyte
  SDRAM 0        REV 01   710-001196   S00013            32 Mbytes
  SDRAM 1        REV 01   710-001196   S00032            32 Mbytes
  PIC 0          REV 03   750-001900   S/N AA9626        1x STM-16 SDH, SMIR
  PIC 1          REV 01   710-002381   S/N AD3633        2x G/E, 1000 BASE-SX
FPC 2                                                    FPC Type OC192
  CPU            REV 03   710-001217   S/N AB3329
  SSRAM          REV 01   710-000077   S/N 306466        1 Mbyte
  SDRAM 0                 710-001196                     32 Mbytes
  SDRAM 1                 710-001196                     32 Mbytes
  SSRAM          REV 01   710-000077   S/N 303919        1 Mbyte
  SDRAM 0                 710-001196                     32 Mbytes
  SDRAM 1                 710-001196                     32 Mbytes
  SSRAM          REV 01   710-000077   S/N 306476        1 Mbyte
  SDRAM 0                                                32 Mbytes
  SDRAM 1                                                32 Mbytes
  SSRAM          REV 01   710-000077   S/N 306409        1 Mbyte
  SDRAM 0                                                32 Mbytes
  SDRAM 1                                                32 Mbytes
  PIC 0          REV 01                                  1x OC-192 SM SR-2
```

show chassis mac-addresses

Syntax show chassis mac-addresses

Description Display the MAC addresses for the router chassis.

Required Privilege Level view

Output Fields Public base address—Base address of the MAC addresses allocated to this router.

Public count—Number of allocated public addresses.

Private base address—Base address of the private MAC addresses allocated to this router.

Private count—Number of allocated private addresses.

Sample Output

```
user@host> show chassis mac-addresses
MAC address information
  Public base address    0:90:69:0:4:0
  Public count           1008
  Private base address   0:90:69:0:7:f0
  Private count          16
```

show chassis routing-engine

Syntax show chassis routing-engine <*slot* >

Description Display information about the Routing Engine.

Options none—Display information about all Routing Engines.

slot—(Optional, for routers with multiple Routing Engines) Display information for an individual Routing Engine. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields Slot—(For routers that support multiple Routing Engines) Slot number.

Current state—(For routers that support multiple Routing Engines) Current state of Routing Engine. It can be Master, Backup, or Disabled.

Election priority—(For routers that support multiple Routing Engines) Election priority for the Routing Engine. It can be Master or Backup.

Temperature—Temperature of the air flowing past the Routing Engine.

DRAM—Total DRAM available to the Routing Engine's processor.

CPU utilization—Information about the Routing Engine's CPU utilization:

- User—Percentage of CPU time being used by user processes.
- Background—Percentage of CPU time being used by background processes.
- Kernel—Percentage of CPU time being used by kernel processes.
- Interrupt—Percentage of CPU time being used by interrupts.
- Idle—Percentage of CPU time that is idle.

Serial ID—(For routers that support multiple Routing Engines) Identification number of Routing Engine in this slot.

Start time—Time at which the Routing Engine started running.

Uptime—How long the Routing Engine has been running.

Load averages—Routing Engine load averages for the last 1, 5, and 15 minutes.

Sample Output: show chassis routing-engine (for M5 router)

```
user@m5-host> show chassis routing-engine
Routing Engine status:
   Temperature                  29 Centigrade
   DRAM                        253 Mbytes
CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
   Start time                      2000-08-28 14:01:06 PDT
   Uptime                          4 hours, 40 minutes, 5 seconds
   Load averages:                  1 minute 5 minute 15 minute
                                      0.00      0.00      0.00
```

Sample Output: show chassis routing-engine (for M10 router)

```
user@m10-host> show chassis routing-engine
Routing Engine status:
    Temperature                 29 degrees C / 84 degrees F
    DRAM                       768 Mbytes
    CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
    Start time                     2000-08-23 19:53:31 PDT
    Uptime                         4 days, 21 hours, 53 minutes, 38 seconds
    Load averages:                 1 minute   5 minute  15 minute
                                       0.12       0.03       0.01
```

Sample Output: show chassis routing-engine (for M20 router)

```
user@m20-host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
    Temperature                 37 Centigrade
    DRAM                       765 Mbytes
    CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                     100 percent
    Serial ID                      98000004f8f27501
    Start time                     2000-03-01 06:27:42 PST
    Uptime                         5 hours, 3 minutes, 35 seconds
    Load averages:                 1 minute   5 minute  15 minute
                                     0.00       0.00       0.00
  Slot 1:                          Empty
```

Sample Output: show chassis routing-engine (for M40 router)

```
user@m40-host> show chassis routing-engine
Routing Engine status:
    Temperature                 48 Centigrade
    DRAM                       253 Mbytes
    CPU utilization:
      User                      51 percent
      Background                 0 percent
      Kernel                    47 percent
      Interrupt                  1 percent
      Idle                       0 percent
    Start time                     2000-03-01 06:57:08 PST
    Uptime                         4 hours, 45 minutes, 59 seconds
    Load averages:                 1 minute   5 minute  15 minute
                                     1.62       1.62       1.54
```

Sample Output: show chassis routing-engine (for M40e router)

```
user@m40E-host> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
    Temperature                 31 degrees C / 87 degrees F
    DRAM                       768 Mbytes
    CPU utilization:
      User                       0 percent
      Background                 0 percent
      Kernel                     0 percent
      Interrupt                  0 percent
      Idle                      99 percent
    Serial ID                      3e000007c8176601
    Start time                     2001-10-17 16:56:14 PDT
    Uptime                         7 days, 22 hours, 1 minute, 22 seconds
    Load averages:                 1 minute   5 minute  15 minute
                                       0.11       0.05       0.01
Routing Engine status:
  Slot 1:
    Current state                  Empty
```

Sample Output: show chassis routing-engine (for M160 router)

```
user@m160-host> show chassis routing-engine
Routing Engine status
 Slot 0
  Current state: Master
  Election priority: Master
  Temperature                41 C / 105 degrees F
  DRAM                       765 Mbytes
  CPU utilization
    User                     0 percent
    Background               0 percent
    Kernel                   0 percent
    Interrupt                0 percent
    Idle                   100 percent
  Serial ID                    39000004f8bdec01
  Start time                    2000-01-04 22:02:58 UTC
  Uptime                        14 hours, 45 minutes, 40 seconds
  Load averages                 1 minute   5 minute  15 minute
                                     0.05       0.04       0.01
 Slot 1
  Current state   Backup
  Election priority   Backup (default)
  Temperature                 41 C / 105 degrees F
  DRAM                       765 Mbytes
  CPU utilization
    User                     0 percent
    Background               0 percent
    Kernel                   0 percent
    Interrupt                2 percent
    Idle                    98 percent
  Serial ID                     f2000004f903a801
  Start time                     2000-01-04 01:28:02 UTC
  Uptime                         20 hours, 38 minutes, 1 seconds
```

show chassis scb

Syntax show chassis scb

Description (M40 router only) Display status information about the SCB.

Required Privilege Level view

Output Fields Temperature—Temperature of the air passing by the SCB, in degrees Centigrade.

CPU utilization—Total percentage of CPU being used by the SCB's processor.

Interrupt utilization—Of the total CPU being used by the SCB's processor, the percentage being used for interrupts.

Heap utilization—Percentage of heap space being used by the SCB's processor.

Buffer utilization—Percentage of buffer space being used by the SCB's processor.

DRAM—Total DRAM available to the SCB's processor.

Start time—Time when the SCB started running.

Uptime—How long the SCB has been running.

Internet Processor memory—Information about the memory of the Internet Processor ASIC on the SCB:

- IP routes—Number of IP routes known to the Internet Processor.
- MPLS routes—Number of MPLS routes known to the Internet Processor.
- SRAM banks enabled—Which SRAM banks are enabled.
- SRAM size—Size of SCB SRAM in bytes.
- SRAM used—Amount of SRAM used, in bytes.
- SRAM utilization—Percentage of SRAM used.

Sample Output

```
user@host> show chassis scb
SCB status:
  Temperature:                30 Centigrade
  CPU utilization:             5 percent
  Interrupt utilization:       0 percent
  Heap utilization:            0 percent
  Buffer utilization:          2 percent
  DRAM:                       64 Mbytes
  Start time:                    1998-10-28 18:35:46 UTC
  Uptime:                        6 minutes, 16 seconds
  Internet Processor memory:
    IP routes:                16
    MPLS routes:               1
    SRAM banks enabled:      [ 1 1 1 1 ]
    SRAM size:                 4 Mbytes
    SRAM used:               256 bytes
    SRAM utilization:          0 percent
```

show chassis sfm

Syntax show chassis sfm <detail>

Description (M40e and M160 routers only) Display status information about the Switching and Forwarding Modules (SFMs).

Options detail—(Optional) Display detailed information.

Required Privilege Level view

Sample Output Sample Output: show chassis sfm (standard) on page 164
Sample Output: show chassis sfm detail on page 164

Output Fields Slot—Slot number. It can be 0, 1, 2 or 3.

Temp—(Standard output only) Temperature of air passing by the SFM, in degrees Centigrade.

CPU Utilization (%)—(Standard output only) Information about CPU usage.

Total—Total percentage of the CPU being used by the SFM's processor.

Interrupt—Of the total CPU being used by the SFM's processor, the percentage being used for interrupts.

Memory Utilization (%)—(Standard output only) Information about memory usage.

DRAM—Total DRAM available to the SFM's processor, in megabytes.

Heap—Percentage of heap space (dynamic memory) being used by the SFM's processor. If this number exceeds 80%, it might indicate a software problem (memory leak).

Buffer—Percentage of buffer space being used by the SFM's processor for buffering internal messages.

State—Status of the Switching and Forwarding Modules. It can be Online, Offline, or Empty.

SPP Temperature—(Detail output only) Temperature of air passing by the Switch Plane Processor card, in degrees Centigrade.

SPR Temperature—(Detail output only) Temperature of air passing by the Switch Plane Router card, in degrees Centigrade.

Total CPU DRAM—(Detail output only) Total amount of CPU DRAM being used by the SFM's processor.

Total SRAM—(Detail output only) Total amount of SRAM being used by the SFM's processor.

Start time—(Detail output only) Time this SFM became active.

Uptime—(Detail output only) How long the SFM has been up and running.

Sample Output: show chassis sfm (standard)

```
user@m160-host> show chassis sfm
SFM status:
                Temp  CPU Utilization (%)   Memory    Utilization (%)
Slot State       (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online       39      0          0        64         0          6
  1  Online       43      0          0        64         0          6
  2  Empty         0      0          0         0         0          0
  3  Empty         0      0          0         0         0          0
```

Sample Output: show chassis sfm detail

```
user@host> show chassis sfm detail
Slot 0 information:
  State:                          Online
  SPP Temperature:            38 Centigrade
  SPR Temperature:            48 Centigrade
  Total CPU DRAM:             64 Mbytes
  Total SRAM:                  4 Mbytes
  Start time:                     2000-01-13 17:58:11 UTC
  Uptime:                         6 hours, 26 minutes, 10 seconds
Slot 1 information:
  State:                          Online
  SPP Temperature:            42 Centigrade
  SPR Temperature:            50 Centigrade
  Total CPU DRAM:             64 Mbytes
  Total SRAM:                  4 Mbytes
  Start time:                     2000-01-13 06:08:02 UTC
  Uptime:                         18 hours, 16 minutes, 19 seconds
```

show chassis ssb

Syntax show chassis ssb <*slot*>

Description (M20 router only) Display status information about the SSB.

Options none—Display information about all SSB slots.

slot—(Optional) Display information about a specific SSB slot.

Required Privilege Level view

Output Fields Failover—Number of times the mastership has changed.

Slot—Name of slot. It can be 0 or 1.

State—Current state of the SSB in this slot. Can be Master, Backup, or Empty.

Temperature—Temperature of the air passing by the SSB, in degrees Centigrade.

CPU utilization—Total percentage of the CPU being used by the SSB's processor.

Interrupt utilization—Of the total CPU being used by the SSB's processor, the percentage being used for interrupts.

Heap utilization—Percentage of heap space being used by the SSB's processor.

Buffer utilization—Percentage of buffer space being used by the SSB's processor.

DRAM—Total DRAM available to the SSB's processor.

Start time—Time when the SSB started running.

Uptime—How long the SSB has been running.

Sample Output

```
user@m20-host> show chassis ssb
SSB status:
  Failover:                      0 time
  Slot 0:
    State:                         Master
    Temperature:                33 Centigrade
    CPU utilization:             0 percent
    Interrupt utilization:       0 percent
    Heap utilization:            0 percent
    Buffer utilization:          6 percent
    DRAM:                       64 Mbytes
    Start time:                    1999-01-15 22:05:36 UTC
    Uptime:                        21 hours, 21 minutes, 22 seconds
  Slot 1:
    State:                         Backup
```

Chapter 6
T-Series Platform Monitoring and Troubleshooting

From the command-line interface (CLI), you can display information about the T-series platform chassis components using the commands in the show chassis hierarchy. These commands get information from the chassis software process, which runs when the router is up and running.

Table 12 summarizes the CLI commands you can use to monitor the routing node. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 12: Commands for Monitoring the T-series platforms

Task or Information to Monitor	CLI Command
Environmental information	show chassis environment on page 169 show chassis environment cb on page 171 show chassis environment pem on page 172 show chassis environment scg on page 173 show chassis environment sib on page 174
Information about the SPMB	show chassis spmb on page 174 show chassis spmb sibs on page 175
Change status	request chassis scg on page 167 request chassis sib on page 168 request chassis spmb restart slot on page 168

request chassis scg

Syntax request chassis scg <offline | online> slot <*slot number*>

Description (T320 and T640 platforms only) Request to take offline the specified SONET Clock Generator (SCG).

Options Offline—Take the SCG offline.

Online—Take the SCG online.

slot number—SCG slot number. It can be 0 or 1.

Required Privilege Level maintenance

request chassis sib

Syntax request chassis sib <offline | online> slot <*slot number*>

Description (T320 and T640 platforms only) Request to take offline the specified SIB.

Options Offline—Take the SIB offline.

Online—Take the SIB online.

slot number—SIB slot number. For the T640, it can be 0 through 4. For the T320, it can be 0 through 2.

Required Privilege Level maintenance

request chassis spmb restart slot

Syntax request chassis spmb restart slot <*slot number*>

Description (T320 and T640 platforms only) Restart the specified Switch Processor Mezzanine Board (SPMB) on the Control Board (CB).

Options *slot number*—CB slot number. It can be 0 or 1.

Required Privilege Level maintenance

show chassis environment

Syntax show chassis environment

Description Display environmental information about the router chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

Required Privilege Level view

Output Fields Power—Information about the Power Entry Modules. Status can be OK, Testing (during initial power-on), Failed, or Absent.

Temp—Temperature of the air flowing through the chassis.

Fans—Information about the fans and blowers. Status can be OK, Testing (during initial power-on), Failed, or Absent. Measurement indicates if spinning at normal or high speed. (The T-320 Internet router has fans only.)

Misc—Information about other components of the chassis. It could indicate the presence of one or more components.

On the T640 routing node, Misc includes CIP (Connector Interface Panel) and SPMB (Switch Processor Mezzanine Board). OK indicates the item is present.

Sample Output: show chassis environment for T640 routing node

```
user@host> show chassis environment

Class Item                  Status     Measurement
Power PEM 0                 OK         22 degrees C / 71 degrees F
      PEM 1                 OK         22 degrees C / 71 degrees F
Temp  SCG 0                 OK         29 degrees C / 84 degrees F
      SCG 1                 OK         28 degrees C / 82 degrees F
      Host 0                OK         28 degrees C / 82 degrees F
      Host 1                OK         25 degrees C / 77 degrees F
      CB 0                  OK         29 degrees C / 84 degrees F
      CB 1                  OK         32 degrees C / 89 degrees F
      SIB 0                 OK         39 degrees C / 102 degrees F
      SIB 1                 OK         42 degrees C / 107 degrees F
      SIB 2                 OK         42 degrees C / 107 degrees F
      SIB 3                 OK         40 degrees C / 104 degrees F
      SIB 4                 OK         0 degrees C / 32 degrees F
      FPC 0 Top             OK         29 degrees C / 84 degrees F
      FPC 0 Bottom          OK         24 degrees C / 75 degrees F
      FPC 1 Top             OK         28 degrees C / 82 degrees F
      FPC 1 Bottom          OK         25 degrees C / 77 degrees F
      FPC 6 Top             OK         34 degrees C / 93 degrees F
      FPC 6 Bottom          OK         24 degrees C / 75 degrees F
      FPM GBUS              OK         23 degrees C / 73 degrees F
      FPM Display           OK         27 degrees C / 80 degrees F
Fans  Top Left Front fan    OK         Spinning at normal speed
      Top Left Middle fan   OK         Spinning at normal speed
      Top Left Rear fan     OK         Spinning at normal speed
      Top Right Front fan   OK         Spinning at normal speed
      Top Right Middle fan  OK         Spinning at normal speed
      Top Right Rear fan    OK         Spinning at normal speed
      Bottom Left Front fan OK         Spinning at normal speed
      Bottom Left Middle fan OK        Spinning at normal speed
      Bottom Left Rear fan  OK         Spinning at normal speed
      Bottom Right Front fan OK        Spinning at normal speed
      Bottom Right Middle fan OK       Spinning at normal speed
      Bottom Right Rear fan OK         Spinning at normal speed
```

```
      Top Blower            OK        Spinning at normal speed
      Second Blower         OK        Spinning at normal speed
      Middle Blower         OK        Spinning at normal speed
      Fourth Blower         OK        Spinning at normal speed
      Bottom Blower         OK        Spinning at normal speed
Misc  CIP                   OK
      SPMB 0                OK
```

Sample Output: show chassis environment for T320 Internet router

```
user@t320-host> show chassis environment
Class Item                  Status     Measurement
Power PEM 0                 OK
      PEM 1                 Absent
Temp  SCG 0                 OK         32 degrees C / 89 degrees F
      SCG 1                 OK         31 degrees C / 87 degrees F
      Routing Engine 0      OK         32 degrees C / 89 degrees F
      CB 0                  OK         37 degrees C / 98 degrees F
      SIB 0                 OK         35 degrees C / 95 degrees F
      SIB 1                 OK         36 degrees C / 96 degrees F
      SIB 2                 OK         39 degrees C / 102 degrees F
      FPC 0 Top             OK         43 degrees C / 109 degrees F
      FPC 0 Bottom          OK         36 degrees C / 96 degrees F
      FPC 1 Top             OK         46 degrees C / 114 degrees F
      FPC 1 Bottom          OK         41 degrees C / 105 degrees F
      FPC 2 Top             OK         44 degrees C / 111 degrees F
      FPC 2 Bottom          OK         37 degrees C / 98 degrees F
      FPC 3 Top             OK         45 degrees C / 113 degrees F
      FPC 3 Bottom          OK         38 degrees C / 100 degrees F
      FPC 4 Top             OK         44 degrees C / 111 degrees F
      FPC 4 Bottom          OK         36 degrees C / 96 degrees F
      FPC 5 Top             OK         46 degrees C / 114 degrees F
      FPC 5 Bottom          OK         41 degrees C / 105 degrees F
      FPC 6 Top             OK         49 degrees C / 120 degrees F
      FPC 6 Bottom          OK         42 degrees C / 107 degrees F
      FPC 7 Top             OK         44 degrees C / 111 degrees F
      FPC 7 Bottom          OK         36 degrees C / 96 degrees F
      FPM GBUS              OK         32 degrees C / 89 degrees F
      FPM Display           OK         34 degrees C / 93 degrees F
Fans  Top Left Front fan    OK         Spinning at normal speed
      Top Left Middle fan   OK         Spinning at normal speed
      Top Left Rear fan     OK         Spinning at normal speed
      Top Right Front fan   OK         Spinning at normal speed
      Top Right Middle fan  OK         Spinning at normal speed
      Top Right Rear fan    OK         Spinning at normal speed
      Bottom Left Front fan OK         Spinning at normal speed
      Bottom Left Middle fan OK        Spinning at normal speed
      Bottom Left Rear fan  OK         Spinning at normal speed
      Bottom Right Front fan OK        Spinning at normal speed
      Bottom Right Middle fan OK       Spinning at normal speed
      Bottom Right Rear fan OK         Spinning at normal speed
      Rear Tray Top fan     OK         Spinning at normal speed
      Rear Tray Second fan  OK         Spinning at normal speed
      Rear Tray Middle fan  OK         Spinning at normal speed
      Rear Tray Fourth fan  OK         Spinning at normal speed
      Rear Tray Bottom fan  OK         Spinning at normal speed
Misc  CIP                   OK
      SPMB 0                OK
```

show chassis environment cb

Syntax show chassis environment cb <*slot*>

Description (T320 and T640 platforms only) Display environmental information about the Control Board (CB).

Options none—Display environmental information about all CBs.

slot—(Optional) Display environmental information about the specified CB. *slot* can be 0 or 1.

Required Privilege Level view

Sample Output

```
user@host> show chassis environment cb
CB 0 status:
  State                            Online Master
  Temperature                   29 degrees C / 84 degrees F
  Power:
    1.8 V                    1805 mV
    2.5 V                    2501 mV
    3.3 V                    3293 mV
    4.6 V                    4725 mV
    5.0 V                    5032 mV
    12.0 V                  11975 mV
    3.3 V bias               3286 mV
    8.0 V bias               7589 mV
  GBUS Revision                 40
  FPGA Revision                  7
CB 1 status:
  State                            Online Standby
  Temperature                   32 degrees C / 89 degrees F
  Power:
    1.8 V                    1802 mV
    2.5 V                    2482 mV
    3.3 V                    3289 mV
    4.6 V                    4720 mV
    5.0 V                    5001 mV
    12.0 V                  11946 mV
    3.3 V bias               3274 mV
    8.0 V bias               7562 mV
  GBUS Revision                 40
  FPGA Revision                  7
```

show chassis environment pem

Syntax show chassis environment pem <*slot*>

Description (T320 and T640 platforms only) Display environmental information about the Power Entry Module (PEM).

Options none—Display environmental information about both PEMs.

slot—(Optional) Display environmental information about an individual PEM. *slot* can be 0 or 1.

Required Privilege Level view

Output Fields slot—Name of PEM slot.

State—Status of PEM.

Temperature—Temperature of the air flowing past the PEM.

DC Input—Status of the DC input.

DC Output—(T640 only) Status of the DC output.

Voltage—(T640 only) Information about voltage of the PEM.

Current—(T640 only) Information about the current of the PEM.

Power—(T640 only) Information about the power of the PEM.

Load—(T640 only) Information about the load on supply, in percentage of rated current being used.

Sample Output: show chassis environment pem for T640 routing node

```
user@host> show chassis environment pem
PEM 0 status:
  State                           Online
  Temperature                 22 degrees C / 71 degrees F
  DC input: OK
  DC Output:            Voltage   Current       Power    Load
       FPC 0              56875       606          34       4
       FPC 1              57016       525          29       3
       FPC 2                  0         0           0       0
       FPC 3                  0         0           0       0
       FPC 4                  0         0           0       0
       FPC 5                  0         0           0       0
       FPC 6              57158      1581          90      12
       FPC 7                  0         0           0       0
  SCG/CB/SIB              56750      1125          63       5
```

Sample Output: show chassis environment pem for T320 routing Internet router

```
user@t320-host> show chassis environment pem
PEM 0 status:
  State                           Online
  Temperature                     OK
  DC input:                       OK
```

show chassis environment scg

Syntax show chassis environment scg <*slot*>

Description (T320 and T640 platforms only) Display environmental information about the SCG.

Options none—Display environmental information about all SCGs.

slot—(Optional) Display environmental information about the specified SCG. *slot* can be 0 or 1.

Required Privilege Level view

Sample Output

```
user@host> show chassis environment scg
SCG 0 status:
  State                            Online -  Master clock
  Temperature                   29 degrees C / 84 degrees F
  Power:
    GROUND                       0 mV
    3.3 V                     3297 mV
    5.0 V                     5050 mV
    5.6 V                     5682 mV
    1.8 V bias                1787 mV
    3.3 V bias                3277 mV
    5.0 V bias                4984 mV
    8.0 V bias                8400 mV
  GBUS Revision                 40
  FPGA Revision                  1.6
SCG 1 status:
  State                            Online - Standby
  Temperature                   28 degrees C / 82 degrees F
  Power:
    GROUND                       0 mV
    3.3 V                     3317 mV
    5.0 V                     5057 mV
    5.6 V                     5689 mV
    1.8 V bias                1794 mV
    3.3 V bias                3296 mV
    5.0 V bias                4991 mV
    8.0 V bias                8410 mV
  GBUS Revision                 40
  FPGA Revision                  1.6
```

show chassis environment sib

Syntax show chassis environment sib <*slot*>

Description (T320 and T640 platforms only) Display environmental information about the SIB.

Options none—Display environmental information about all SIBs.

slot—(Optional) Display environmental information about the specified SIB.
For the T640, *slot* can be 0 through 4. For the T320, *slot* can be 0 through 2.

Required Privilege Level view

Sample Output

```
user@host> show chassis environment sib 1
SIB 1 status:
  State                           Online
  Temperature                  39 degrees C / 102 degrees F
  Power:
    GROUND                      0 mV
    1.8 V                    1809 mV
    2.5 V                    2478 mV
    3.3 V                    3308 mV
    1.8 V bias               1794 mV
    3.3 V bias               3274 mV
    5.0 V bias               4996 mV
    8.0 V bias               7247 mV
```

show chassis spmb

Syntax show chassis spmb

Description (T320 and T640 platforms only) Display information about the SPMB.

Required Privilege Level view

Sample Output

```
user@host> show chassis spmb
Slot 0 information:
  State                          Online
  Total CPU Utilization             1%
  Interrupt CPU Utilization         0%
  Memory Heap Utilization           0%
  Buffer Utilization               40%
  Start time:                    2001-08-27 14:05:04 PDT
  Uptime:                         46 minutes, 36 seconds
```

show chassis spmb sibs

Syntax show chassis spmb sibs

Description (T320 and T640 platforms only) Display state of the SIBs.

Required Privilege Level view

Output Fields Slot—SIB slot number. It can be 0 through 4. For the T640, slot can be 0 through 4. For the T320, slot can be 0 through 2.

State—State of the SIB. Possible values are Online, Offline, Spare, Empty, Fault, and Check.

Sample Output: show chassis spmb sibs (for T320)

```
user@t320-host> show chassis spmb sibs


Slot  State
 0    Spare
 1    Online
 2    Online
```

Sample Output: show chassis spmb sibs (for T640)

```
user@t640-host> show chassis spmb sibs


Slot  State
 0    Online
 1    Online
 2    Empty
 3    Online
 4    Offline
```

Chapter 7

Accounting Options Monitoring and Troubleshooting

Table 13 summarizes the command-line interface (CLI) commands you can use to monitor accounting options. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 13: Commands for Monitoring and Managing Accounting Options

Task or Information to Monitor	Command
Display accounting statistics	show accounting profile on page 178
Display records	show accounting records on page 182

show accounting profile

Syntax show accounting profile *profile-name*

Description Display accounting profile information.

Options *profile-name*—Name of accounting profile.

Privilege Level view

Output Fields Profile—Name of accounting profile.

Sampling interval—Configured interval, in minutes, for statistic collection. The minimum interval is one minute. The range is 1 through 2880 minutes. The default is 30 minutes.

Profile Usage Count—Number of items configured for collecting accounting statistics.

File information—Information about the accounting profile log, including:

- File—Name of accounting profile log. If no name is explicitly provided, the name of the accounting profile is used. All statistics files are placed in the /var/log directory.
- maximum size—Configured size. When size is exceeded, the log file closes and a new log file opens. The default size is 128 KB.
- maximum number—Configured maximum number of log files. It can be between 3 and 1000. The default number is 10.
- bytes written—Number of bytes written to the log file.

Transfer Interval—Length of time (in minutes) the file remains open, receiving statistics before it is closed, transferred, and rotated. The range is 15 minutes to 48 hours (2880 minutes). The default is 30 minutes. When either the time or the file size is exceeded, the file is closed and a new one opened, whether or not there is a transfer site specified.

Next Scheduled Transfer—Time at which next transfer occurs.

Column Labels—Names of sampled statistics. List varies depending on configuration.

- profile-layout—List of data fields reported, in the order they appear in the output.
- epoch-timestamp—Number of seconds since the Epoch.
- interfaces—(For interface, filter, and destination class profiles) Name of the interfaces on which the filter is applied.
- filter-name—(For filter profiles) Name of the filter.
- counter-name—(For filter profiles) Name of the counter.
- packet-count—(For filter and destination class profiles) Number of packets for the counter.
- byte-count—(For filter and destination class profiles) Number of bytes for the counter.
- input-bytes—(For interface profiles) Input bytes.

- input-errors—(For interface profiles) Generic input error packets.
- input-multicast—(For interface profiles) Input packets arriving by multicast.
- input-packets—(For interface profiles) Input packets.
- input-unicast—(For interface profiles) Input unicast packets.
- output-bytes—(For interface profiles) Output bytes.
- output-errors—(For interface profiles) Generic output error packets.
- output-multicast—(For interface profiles) Output packets sent by multicast.
- output-packets—(For interface profiles) Output packets.
- output-unicast—(For interface profiles) Output unicast packets.
- no-proto—(For interface profiles) Packets for unsupported protocol.
- snmp-index—(For interface profiles) SNMP index.
- destination-class-name—(For destination class profiles) Configured destination class name.
- host name—(For Routing Engine profiles) Host name for the router.
- date-*yyyymmdd*—(For Routing Engine profiles) Date in YYYMMDD format.
- timeofday-*hhmmss*—(For Routing Engine profiles) Time of day in HHMMSS format.
- uptime—(For Routing Engine profiles) Time since last reboot, in seconds.
- cpu1min—(For Routing Engine profiles) Average system load over the last 1 minute.
- cpu5min—(For Routing Engine profiles) Average system load over the last 5 minutes.
- cpu15min—(For Routing Engine profiles) Average system load over the last 15 minutes.

Interface Name—Name of interfaces configured for this accounting profile.

Filter Name—Name of filter configured for this accounting profile.

Routing Engine stats—Routing engine accounting profile.

Destination class Name—Name of destination class configured for this accounting profile.

Next Scheduled Collection—Time for next collection of statistics for the named interface.

Sample output: show accounting profile

```
user@host> show accounting profile if_prof
Profile if_prof
Sampling interval: 1 minute(s),  Profile Usage Count: 2
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes
written 2196
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:45
Column Labels:
  profile-layout
  epoch-timestamp
  interface-name
  snmp-index
  input-bytes
  output-bytes
  input-packets
  output-packets
  input-unicast
  output-unicast
  input-multicast
  output-multicast
  no-proto
  input-errors
  output-errors



Interface Name              Next Scheduled Collection
fxp0.0                      2001-06-18-18:00:30
fxp0                        2001-06-18-18:01:00
```

Sample output: show accounting profile (for filter)

```
user@host> show accounting profile filter_profile
Profile filter_profile
Sampling interval: 1 minute(s),  Profile Usage Count: 0
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes
written 822
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  interfaces
  filter-name
  counter-name
  packet-count
  byte-count



Filter Name                 Next Scheduled Collection
myfiltero                         2001-06-03-04:32:59
```

Sample output: show accounting profile (for destination class)

```
user@host> show accounting profile dcu1
Profile dcu1
Sampling interval: 1 minute(s),  Profile Usage Count: 0
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes
written 901
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  interface-name
  destination-class-name
  packet-count
  byte-count



Interface Name              Next Scheduled Collection
so-0/3/3                    2001-06-03-04:34:00
```

Sample output: show accounting profile (for Routing Engine)

```
user@host> show accounting profile rep1
Profile rep1
Sampling interval: 1 minute(s),  Profile Usage Count: 1
File accounting_profile_stats: maximum size 1048576, maximum number 5, bytes
written 901
Transfer Interval: 15 minute(s),  Next Scheduled Transfer: 2001-06-17-18:00:46
Column Labels:
  profile-layout
  epoch-timestamp
  hostname
  date-yyyymmdd
  timeofday-hhmmss
  uptime
  cpu1min
  cpu5min
  cpu15min



Interface Name              Next Scheduled Collection
routing-engine-stats        2001-06-18-18:02:31
```

show accounting records

Syntax show accounting records *profile-name* <since *time*> <utc_timestamp>

Description Display accounting records for the specified accounting profile.

Options *profile-name*—Name of accounting profile.

since *time*—(Optional) Display accounting statistics since the specified time (YYYY-MM-DD-HH:MM:SS).

utc_timestamp—(Optional) Display timestamp in UTC format.

Privilege Level view

Output Fields Timestamp—Date and time of sample.

Interface Name—(For interface profiles) Name and SNMP index of interface for accounting profile.

Filter Name—(For filter profiles) Name of filter.

Interfaces—(For filter profiles) Name of interfaces for accounting profile.

Counter Name—(For filter profiles) Name of counter.

Destination Class—(For Destination Class profiles) Name of Destination Class.

Input Bytes—(For interface profiles) Input bytes.

Output Bytes—(For interface profiles) Output bytes.

Input Packets—(For interface profiles) Input packets.

Output Packets—(For interface profiles) Output packets.

Input Unicast Packets —(For interface profiles) Input unicast packets.

Output Unicast Packets—(For interface profiles) Output unicast packets.

Input Multicast Packets—(For interface profiles) Input packets arriving by multicast.

Output Multicast Packets—(For interface profiles) Output packets sent by multicast.

Unsupported Protocol Packets—(For interface profiles) Packets for unsupported protocol.

Input Errors—(For interface profiles) Generic input error packets.

Output Errors—(For interface profiles) Generic output error packets.

Host Name—(For Routing Engine profiles) Host name for the router.

Date—(For Routing Engine profiles) Date in YYYMMDD format.

Time of Day—(For Routing Engine profiles) Time of day in HHMMSS format.

Uptime—(For Routing Engine profiles) Time since last reboot, in seconds.

Average CPU Load (1 min)—(For Routing Engine profiles) Average system load over the last 1 minute.

Average CPU Load (5 min)—(For Routing Engine profiles) Average system load over the last 5 minutes.

Average CPU Load (15 min)—(For Routing Engine profiles) Average system load over the last 15 minutes.

Sample Output

```
user@host> show accounting records if_prof
Timestamp: 2000-10-03-00:30:41, Interface Name: fxp0 (SNMP Index 1)
             32663634   Input Bytes
              3487515   Output Bytes
               158000   Input Packets
                33296   Output Packets
               158000   Input Unicast Packets
                33296   Output Unicast Packets
                    0   Input Multicast Packets
                    0   Output Multicast Packets
                    0   Unsupported Protocol Packets
                    0   Input Errors
                    0   Output Errors
```

Sample Output: show accounting records (with utc-timestamp)

```
user@host> show accounting records if_prof utc_timestamp
Timestamp: 2001-06-18-18:01:00, Interface Name: fxp0 (SNMP Index 1)
             32663634   Input Bytes
              3487515   Output Bytes
               158000   Input Packets
                33296   Output Packets
               158000   Input Unicast Packets
                33296   Output Unicast Packets
                    0   Input Multicast Packets
                    0   Output Multicast Packets
                    0   Unsupported Protocol Packets
                    0   Input Errors
                    0   Output Errors
```

Sample Output: show accounting records (with since time)

```
user@host> show accounting records if_prof since 2000-10-03-00:10:41
Timestamp: 2000-10-03-00:30:41, Interface Name: fxp0 (SNMP Index 1)
             32663634   Input Bytes
              3487515   Output Bytes
               158000   Input Packets
                33296   Output Packets
               158000   Input Unicast Packets
                33296   Output Unicast Packets
                    0   Input Multicast Packets
                    0   Output Multicast Packets
                    0   Unsupported Protocol Packets
                    0   Input Errors
                    0   Output Errors
```

Sample Output: show accounting records (for filter profile)

```
user@host> show accounting records filter_profile
Timestamp: 2000-10-03-00:30:41, Filter Name: ap_filter, Interfaces: fxp0.0
    Counter Name: c1
                  2440   Packets
                223509   Bytes
```

Sample Output: show accounting records (for destination class profile)

```
user@host> show accounting records dcu1
Timestamp: 2000-10-03-00:30:41, Interface: so-2/0/0.0, Destination Class: gold
                     0   Packets
                     0   Bytes
```

Sample Output: show accounting records (for Routing Engine profile)

```
user@host> show accounting records rep1
Timestamp: 2000-10-03-00:30:41
Host Name:                      andro
Date:                           20010618
Time of Day:                    183130
Uptime:                         88260
Average CPU Load (1 min):       0.000000
Average CPU Load (5 min):       0.000000
Average CPU Load (15 min):      0.000000
```

Part 3
Interfaces

- Static Router Interfaces Monitoring and Troubleshooting on page 187
- ATM Interfaces Monitoring and Troubleshooting on page 209
- Channelized DS-3 to DS-0 Interfaces Monitoring and Troubleshooting on page 225
- Channelized DS-3 to DS-1 Interfaces Monitoring and Troubleshooting on page 239
- Channelized E-1 Interfaces Monitoring and Troubleshooting on page 253
- Channelized OC-12 Interfaces Monitoring and Troubleshooting on page 269
- Channelized STM-1 Interfaces Monitoring and Troubleshooting on page 283
- E-1 Interfaces Monitoring and Troubleshooting on page 299
- E-3 Interfaces Monitoring and Troubleshooting on page 313
- ES Interfaces Monitoring and Troubleshooting on page 325
- Fast Ethernet Interfaces Monitoring and Troubleshooting on page 335
- Gigabit Ethernet Interfaces Monitoring and Troubleshooting on page 347
- GRE Interfaces Monitoring and Troubleshooting on page 361
- IP-IP Interfaces Monitoring and Troubleshooting on page 371
- Multilink Services Interfaces Monitoring and Troubleshooting on page 381
- SONET/SDH Interfaces Monitoring and Troubleshooting on page 393
- T1 Interfaces Monitoring and Troubleshooting on page 409
- T3 Interfaces Monitoring and Troubleshooting on page 421
- Virtual Tunnel Interfaces Monitoring and Troubleshooting on page 435
- Real-Time Router Interface Monitoring and Troubleshooting on page 445
- Monitoring and Troubleshooting using BERT Tests on page 459
- VRRP Router Interfaces Monitoring and Troubleshooting on page 465

Chapter 8

Static Router Interfaces Monitoring and Troubleshooting

Table 14 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot static router interfaces. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 14: Commands for Monitoring Interfaces

Task or Information to Monitor	CLI Command
General interface information.	show interfaces on page 194
Summary interface information.	show interfaces terse on page 206
Configured interface descriptions.	show interfaces descriptions on page 195
Routing information for each interface.	show interfaces routing on page 200
Media-specific interface information.	show interfaces media on page 199
Interface statistics.	show interfaces statistics on page 205
Clear (zero) interface statistics.	clear interfaces statistics on page 193
Filters or policers installed.	show interfaces (filters | policers) on page 196

Wildcard Characters in Interface Names

With the show interfaces and clear interfaces commands, you can use wildcard characters when specifying interface names (in the *interface-name* option) to refer to groups of interface names without having to type each name individually. Table 15 lists the available wildcard characters.

You must enclose all wildcard characters except the asterisk in quotation marks (“ ”).

Table 15: Wildcard Characters for Specifying Interface Names

Wildcard Character	Description
* (asterisk)	Match any string of characters in that position in the interface name. For example, so* matches all SONET/SDH interfaces.
"[*character* <*character*...>]"	Match one or more individual characters in that position in the interface name. For example, "so-[03]"* matches all SONET/SDH interfaces in slots 0 and 3.
"[!*character* <*character*...>]"	Match all characters except the ones included in the brackets. For example, "so-[!03]"* matches all SONET/SDH interfaces except for those in slots 0 and 3.
"[*character1-character2*]"	Match a range of characters. For example, so-"[0-3]"* matches all SONET/SDH interfaces in slots 0, 1, 2, and 3.
"[!*character1-character2*]"	Match all characters that are not in the specified range of characters. For example, so-"[!0-3]"* matches all SONET/SDH interfaces in slots 4, 5, 6, and 7.

Examples: Use Wildcard Characters in Interface Names

List only SONET/SDH interfaces:

```
user@host> show interfaces terse so*
Interface       Admin Link Proto Local                 Remote
so-1/0/0        up    up
so-1/0/0.0      up    up   inet  192.168.8.192     --> 192.168.2.250
                           iso
so-1/1/0        up    down
so-1/1/0.0      up    down inet  192.168.8.109/30
                           iso
so-1/1/1        up    up
so-1/1/1.0      up    down inet  192.168.8.113/30
                           iso
                           mpls
...
so-1/3/3        up    down
so-5/0/0        up    up
so-5/0/0.0      up    down inet  192.168.8.208     --> 192.168.8.209
                           iso
                           mpls
so-5/0/1        down  down
so-5/0/2        up    down
...
so-5/1/3        up    down
```

List only SONET/SDH interfaces in slot 1:

```
user@host> show interfaces terse so-"[1]"*
Interface       Admin Link Proto Local                 Remote
so-1/0/0        up    up
so-1/0/0.0      up    up   inet  192.168.2.125     --> 192.168.2.250
                           iso
so-1/1/0        up    down
so-1/1/0.0      up    down inet  192.168.2.150/30
                           iso
so-1/1/1        up    up
so-1/1/1.0      up    down inet  192.168.2.175/30
                           iso
                           mpls
so-1/1/2        up    up
```

```
so-1/1/2.0      up    down inet  192.168.2.110/30
                           iso
so-1/1/3        up    up
so-1/1/3.0      up    up   inet  192.168.2.120/30
                           iso
so-1/2/0        up    down
so-1/2/1        up    down
so-1/2/2        up    down
so-1/2/3        up    down
so-1/3/0        up    down
so-1/3/1        up    down
so-1/3/2        up    down
so-1/3/3        up    down
```

List all interfaces whose type starts with “s” or “g”:

```
user@host> show interfaces terse “[sg]”*
Interface       Admin Link Proto Local                 Remote
so-1/0/0        up    up
so-1/0/0.0      up    up   inet  192.168.2.125     --> 192.168.2.250
                           iso
so-1/1/0        up    down
so-1/1/0.0      up    down inet  192.168.2.150/30
                           iso
so-1/1/1        up    up
so-1/1/1.0      up    down inet  192.168.2.175/30
                           iso
                           mpls
so-1/1/2        up    up
...
so-5/1/3        up    down
gre             up    up
```

Common Output Fields

Many of the fields in the output of the various show interfaces commands are the same. This section explains the output of these common fields:

- Enabled Field on page 190
- Device Flags Field on page 190
- Link Flags Field on page 191
- Interface Flags Field on page 191
- Logical Interface Flags Field on page 192
- Family Flags Field on page 192
- Address Flags Field on page 192

Enabled Field

The Enabled field shows the state of the interface. It can be one of the following:

- Administratively down, Physical link is Down—The interface is turned off, and the physical link is unhealthy and could not pass packets if the interface were turned on.
- Administratively down, Physical link is Up—The interface is turned off, and the physical link is healthy and could pass packets if the interface were turned on.
- Enabled, Physical link is Down—The interface is turned on, but the physical link is unhealthy and cannot pass packets.
- Enabled, Physical link is Up—The interface is turned on, and the physical link is healthy and can pass packets.

Device Flags Field

The Device Flags field shows information about the physical device. It can be one or more of the following:

- Down—Device has been administratively disabled.
- Hear-Own-Xmit—Device will hear its own transmissions.
- Link-Layer-Down—The link layer protocol has failed to successfully connect with the remote end point.
- Loopback—Device is in physical loopback.
- Loop-Detected—The link layer has received frames that it sent and suspects a physical loopback.
- No-Carrier—Where the media supports carrier recognition, this indicates that no carrier is currently seen.
- No-Multicast—Device does not support multicast traffic.
- Present—Device is physically present and recognized.
- Promiscuous—Device is in promiscuous mode and will see frames addressed to all physical addresses on the media.
- Quench—Device is quenched because it overran its output buffer.
- Recv-All-Multicasts—No multicast filtering (multicast promiscuous).
- Running—Device is active and enabled.

Link Flags Field

The Link flags field shows information about the physical link. It can be one or more of the following:

- Give-Up—Link protocol will not continue to retry to connect after repeated failures.
- Keepalives—Link protocol keepalives are enabled.
- Loose-LCP—PPP will not use LCP to indicate whether the link protocol is up.
- Loose-LMI—Frame Relay will not use LMI to indicate whether the link protocol is up.
- Loose-NCP—PPP will not use NCP to indicate whether the device is up.
- No-Keepalives—Link protocol keepalives are disabled.

Interface Flags Field

The Interface flags field shows information about the interface. It can be one or more of the following:

- Admin-Test—Interface is in test mode, which means that some sanity checking, such as loop detection, is disabled.
- Disabled—Interface is administratively disabled.
- Hardware-Down—Interface is nonfunctional or incorrectly connected.
- Link-Layer-Down—Interface keepalives have indicated that the link is incomplete.
- No-Multicast—Interface does not support multicast traffic.
- Point-To-Point—Interface is point to point.
- Promiscuous—Interface is in promiscuous mode and will see frames addressed to all physical addresses.
- Recv-All-Multicasts—No multicast filtering (multicast promiscuous).
- SNMP-Traps—SNMP traps are enabled.
- Up—Interface is enabled and operational.

Logical Interface Flags Field

The Logical interface flags field shows information about the logical interface. It can be one or more of the following:

- Disabled—Interface is administratively disabled.
- Hardware-Down—Interface protocol initialization failed to complete successfully.
- Point-To-Point—Interface is point to point.
- SNMP-Traps—SNMP traps are enabled.
- Up—Interface is enabled and operational.

Family Flags Field

The Family flags field shows information about the protocol family on the logical interface. It can be one or more of the following:

- Dest-route-down—The software noticed that the link is down and has stopped forwarding the link's interface routes.
- Down—Protocol is inactive.
- Is-Primary—The interface is the primary one for the protocol.
- No-Redirects—Protocol redirects are disabled.
- Primary—Interface can be considered for selection as the primary family address.
- Protocol-Down—Protocol failed to negotiate correctly.
- Up—Protocol is configured and operational.

Address Flags Field

The Address flags field shows information about the addresses configured for the protocol family. It can be one or more of the following:

- Dest-route-down—RPD noticed that the link was down and changed the interface routes to be nonforwarding.
- Is-Default—This address is the default address of the router. That is, it is the address sourced by SNMP, ping, traceroute, and other network utilities.
- Is-Preferred—This address is the default local address used for packets sourced by the local router to destinations on the subnet.

- Is-Primary—This address is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface.
- Preferred—This address is a candidate to become the preferred address.
- Primary—This address is a candidate to become the primary address.

clear interfaces statistics

Syntax clear interfaces statistics (*interface-name* | all)

Description Zero interface statistics.

Options all—Zero statistics on all interfaces.

interface-name—Name of the interface.

Required Privilege Level clear

See Also monitor interface on page 449, show interfaces statistics on page 205

show interfaces

Syntax show interfaces <detail | extensive> <*interface-name*> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

Description Display static status information about router interfaces.

Options none—Display information about all interfaces.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

interface-name—(Optional) Name of an interface.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

Required Privilege Level view

Sample Output
show interfaces descriptions on page 195
show interfaces (for ATM Interfaces) on page 209
show interfaces (for Channelized DS-3 to DS-0 Interfaces) on page 225
show interfaces (for Channelized DS-3 to DS-1 Interfaces) on page 239
show interfaces (for Channelized E1 Interfaces) on page 253
show interfaces (for Channelized OC-12 Interfaces) on page 269
show interfaces (for Channelized STM-1 Interfaces) on page 283
show interfaces (for E1 Interfaces) on page 299
show interfaces (for E3 Interfaces) on page 313
show interfaces (for ES Interfaces) on page 325
show interfaces (for Fast Ethernet Interfaces) on page 335
show interfaces (for Gigabit Ethernet Interfaces) on page 347
show interfaces (for GRE Interfaces) on page 361
show interfaces (for IP-IP Interfaces) on page 371
show interfaces (for Multilink Services Interfaces) on page 381
show interfaces (for SONET/SDH Interfaces) on page 393
show interfaces (for T1 Interfaces) on page 409
show interfaces (for T3 Interfaces) on page 421
show interfaces (for VT Interfaces) on page 435

show interfaces descriptions

Syntax show interfaces descriptions

Description Display the interface descriptions that have been configured with the description statement at the [edit interfaces *interface-name*] or [edit interfaces *interface-name* unit *unit-number*] hierarchy level.

Required Privilege Level view

Output Fields Interface—Name of interface.

Admin—Whether the interface is turned on (up) or off (down).

Link—Link state. It can be up or down.

Description—Configured description for the interface.

Sample Output

```
user@host> show interfaces descriptions
Interface       Admin Link Description
t1-0/0/0:0      up    down juniper
t1-0/0/0:1      up    down cisco
t1-0/0/0:2      up    down bombay
t1-0/0/0:3      up    down sanjose
t1-0/0/0:4      up    down Atlanta
t1-0/0/0:5      up    down Boston
so-0/1/0        up    down sonet
t3-0/3/0:0      up    up   malayeshia
fe-1/2/1        up    up   ixia_a
fe-1/2/2        up    down ixia_b
ge-1/3/0        up    up   not_connected
fxp0            up    up   fxp0
```

show interfaces (filters | policers)

Syntax show interfaces (filters | policers) < *interface-name* >

Description Display all firewall filters or policers that are installed on each interface in a system.

Options none—Display filter or policer information about all interfaces.

interface-name—(Optional) Display filter or policer information about a particular interface.

Required Privilege Level view

Output Fields Interface—Name of the interface.

Admin—Whether the interface is turned on (up) or off (down).

Link—Link state. It can be up or down.

Proto—Protocol configured on the interface.

Input Policer—(For policer display only) Policer to be evaluated when packets are received on the interface. It has the format *type-fpc/pic/port*-in-policer.

Output Policer—(For policer display only) Policer to be evaluated when packets are transmitted on the interface. It has the format *type-fpc/pic/port*-out-policer.

Input Filter—(For filter display only) Name of the firewall filter to be evaluated when packets are received on the interface.

Output Filter—(For filter display only) Name of the firewall filter to be evaluated when packets are transmitted on the interface.

Sample Output: show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    up
ge-0/0/0.0      up    up   inet
                           iso
gr-0/3/0        up    up
ip-0/3/0        up    up
mt-0/3/0        up    up
pd-0/3/0        up    up
pe-0/3/0        up    up
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
at-1/2/0        up    up
at-1/2/0.0      up    up   inet
                           iso
at-1/2/1        up    down
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  so-2/0/0.0-in-policer so-2/0/0.0-out-policer
                           iso
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
at-2/2/0        up    down
so-3/0/0        up    down
so-4/0/0        up    down
so-5/0/0        up    down
at-6/0/0        up    up
fxp0            up    up
fxp0.0          up    up   inet
fxp1            up    up
fxp1.0          up    up   tnp
gre             up    up
ipip            up    up
lo0             up    up
lo0.0           up    up   inet
lsi             up    up
mtun            up    up
pimd            up    up
pime            up    up
tap             up    up
```

Sample Output: show interfaces policers (interface-name)

```
user@host> show interfaces policers so-2/1/0
Interface       Admin Link Proto Input Policer         Output Policer
so-2/1/0        up    down
so-2/1/0.0      up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                           iso
                           inet6
```

Sample Output: 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
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
at-1/2/0        up    up
at-1/2/0.0      up    up   inet
                           iso
at-1/2/1        up    down
so-2/0/0        up    up
so-2/0/0.0      up    up   inet v4top                outfilter
                           iso
so-2/1/0        up    down
so-2/1/0.0      up    down inet  v5top                 outfilter
                           iso
                           inet6 v6in                v6out
at-2/2/0        up    down
so-3/0/0        up    down
so-4/0/0        up    down
so-5/0/0        up    down
at-6/0/0        up    up
fxp0            up    up
fxp0.0          up    up   inet
fxp1            up    up
fxp1.0          up    up   tnp
gre             up    up
ipip            up    up
lo0             up    up
lo0.0           up    up   inet
lsi             up    up
mtun            up    up
pimd            up    up
pime            up    up
tap             up    up
```

Sample Output: show interfaces filters (interface-name)

```
user@host> show interfaces filters so-2/1/0
Interface       Admin Link Proto Input Filter         Output Filter
so-2/1/0        up    down
so-2/1/0.0      up    down inet  goop                  outfilter
                           iso
                           inet6 v6in                v6out
```

show interfaces media

Syntax show interfaces media <*interface-name*> <detail>

Description Display media-specific information about network interfaces.

Options detail—(Optional) Display detailed interface information.

interface-name—(Optional) Name of an interface.

Required Privilege Level view

Sample Output
show interfaces media (for ATM Interfaces) on page 222
show interfaces media (for Channelized DS-3 to DS-0 Interfaces) on page 237
show interfaces media for Channelized DS-3 to DS-1 Interfaces) on page 250
show interfaces media (for Channelized E1 interfaces) on page 266
show interfaces media (for Channelized OC-12 Interfaces) on page 281
show interfaces media (for Channelized STM-1 Interfaces) on page 296
show interfaces media (for E1 interfaces) on page 311
show interfaces media (for E3 interfaces) on page 322
show interfaces media (for ES interfaces) on page 332
show interfaces media (for Fast Ethernet interfaces) on page 345
show interfaces media (for Gigabit Ethernet interfaces) on page 359
show interfaces media (for GRE Interfaces) on page 369
show interfaces media (for IP-IP interfaces) on page 378
show interfaces media (for Multilink Services Interfaces) on page 391
show interfaces media (for SDH/SONET interfaces) on page 408
show interfaces media (for T1 Interfaces) on page 419
show interfaces media (for T3 interfaces) on page 432
show interfaces media (for VT interfaces) on page 443

show interfaces routing

Syntax show interfaces routing <*interface-name*> <brief | detail | summary>

Description Have the routing protocol process display its view of the state of the router's interfaces.

Use this command for performing router diagnostics only, when you are checking whether the routing protocols and the JUNOS software disagree about the state of an interface.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

interface-name—(Optional) Name of a specific interface.

summary—(Optional) Summarize the interface information.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show interfaces routing brief on page 203
Sample Output: show interfaces routing summary on page 203
Sample Output : show interfaces routing detail on page 204

Options at a Glance Table 16 summarizes which information is included in the output of each show interfaces routing command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 16: Show Interfaces Routing Output Field Summary

Options	Field Description
Brief	Addresses—Protocols and addresses configured on the interface.
Detail	Bandwidth—Speed at which the interface is running.
Detail	Change—Reflects recent changes to the interface.
Detail	Encapsulation—Encapsulation on the interface.
Summary Detail	Index—Interface's index number, which reflects its initialization sequence.
Brief	Interface—Name of the physical interface.
Detail	Link layer—Describes the link layer of the interface.
Summary Detail	Metric—Metric value on the interface.
Detail	Protocol information—information about the configuration of protocols on the interface: ■ Address—Address configured on the interface for the protocol type. ■ State—State of the logical interface (Up or Down) and the type of interface. ■ Change—Reflects recent changes to the interface. ■ Metric—Metric value on the interface for the protocol type. ■ MTU—MTU value of the interface. ■ Local address—On a point-to-point link, the address of the local side of the link. Not used for multicast links. ■ Destination—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address.
Detail	Refcount—Number of references to the interface in the routing software.
Detail	State—State of the logical interface (Up or Down) and the type of interface.
Brief	State—State of the physical interface. It can be Up or Down.
Summary	Status—State of the logical interface (Up or Down) and the type of interface.
Summary	Trans.—Transitions on the interface (number of times the interface has gone from down to up).
Detail	Up/down transitions—Number of times the interface has gone from down to up.

Output Fields

Addresses—(Brief output only) Protocols and addresses configured on the interface.

Interface—(Brief output only) Name of the physical interface.

State—(Brief output only) State of the physical interface. It can be Up or Down.

State—(Detail output only) State of the logical interface (Up or Down) and the type of interface.

Status—(Summary output only) State of the logical interface (Up or Down) and the type of interface.

Trans—(Summary output only) Transitions on the interface (number of times the interface has gone from down to up).

Index—Interface's index number, which reflects its initialization sequence.

Refcount—Number of references to the interface in the routing software.

State—State of the logical interface (Up or Down) and the type of interface.

Change—Reflects recent changes to the interface; value can be one or more of the following:

- Add—The interface was just added.
- Address—The interface's link layer address has changed.
- Delete—The interface is being deleted.
- Encapsulation—The type of encapsulation on the interface has changed.
- Metric—The interface's metric value has changed.
- MTU—The interface's MTU size has changed.
- UpDown—The interface has made an up/down transition.

Metric—Metric value on the interface.

Up/down transitions—Number of times the interface has gone from down to up.

Link layer—Describes the link layer of the interface.

Encapsulation—Encapsulation on the interface.

Bandwidth—Speed at which the interface is running.

Protocol information—information about the configuration of protocols on the interface:

- Address—Address configured on the interface for the protocol type.
- State—State of the logical interface (Up or Down) and the type of interface.
- Change—Reflects recent changes to the interface. It can be one or more of the following:
  - Add—The interface was just added.
  - Address—The interface's address has changed.
  - Broadcast—The interface's broadcast address has changed.
  - Delete—The interface is being deleted.
  - Netmask—The interface's netmask has changed.
  - UpDown—The interface has made an up/down transition.
- Preference—Preference value for the route for this address.
- Metric—Metric value on the interface for the protocol type.
- MTU—MTU value of the interface.

- Local address—On a point-to-point link, the address of the local side of the link. Not used for multicast links.
- Destination—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address.

Sample Output: show interfaces routing brief

```
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
so-5/0/0.0       Up    MPLS  enabled
                       ISO   enabled
                       INET  192.168.2.100
                       INET  enabled
so-1/2/0.0       Down  ISO   enabled
so-1/1/0.0       Up    INET  192.168.2.250
                       INET  enabled
at-1/0/0.6       Up    CCC   enabled
at-1/0/0.5       Up
at-1/0/0.4       Up
at-1/0/0.3       Up    CCC   enabled
at-1/0/0.2       Up    CCC   enabled
at-1/0/0.0       Up    ISO   enabled
                       INET  192.168.90.10
                       INET  enabled
lo0.0            Up    ISO   47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
                       ISO   enabled
                       INET  127.0.0.1
fxp1.0           Up
fxp0.0           Up    INET  192.168.6.90
```

Sample Output: show interfaces routing summary

```
user@host> show interfaces routing summary
14 physical interfaces (12 up)
  11 INET protocol addresses (11 up)
  6 ISO protocol addresses (4 up)
  3 MPLS protocol addresses (3 up)
  3 CCC protocol addresses (3 up)

Interface  Index   Metric   Trans. Status
so-5/0/3.0    15        0        0 Broadcast PointToPoint Multicast
so-5/0/2.0    14        0        0 Up Broadcast PointToPoint Multicast
so-5/0/1.0    13        0        5 Up Broadcast PointToPoint Multicast
so-5/0/0.0    12        0        2 Up Broadcast PointToPoint Multicast
so-1/2/0.0    11        0        0 Broadcast PointToPoint Multicast
so-1/1/0.0    10        0        5 Up Broadcast PointToPoint Multicast
at-1/0/0.6     9        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.5     8        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.4     7        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.3     6        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.2     5        0        0 Up Broadcast PointToPoint Multicast
at-1/0/0.0     4        0        0 Up Broadcast PointToPoint Multicast
lo0.0          3        0        0 Up Broadcast Loopback Multicast
fxp1.0         2        0        1 Up Broadcast Multicast
fxp0.0         1        0        0 Up Broadcast Multicast
```

Sample Output: show interfaces routing detail

```
user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: <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 statistics

Syntax show interfaces statistics <*interface-name*> <detail>

Description Display static interface statistics, such as errors.

Options detail—(Optional) Display detailed information.

interface-name—(Optional) Name of the interface.

Default: detail

Required Privilege Level view

See Also clear interfaces statistics on page 193, monitor interface on page 449

Sample Output show interfaces statistics (for ATM Interfaces) on page 222
show interfaces statistics (for Channelized DS-3 to DS-0 Interfaces) on page 237
show interfaces statistics for Channelized DS-3 to DS-1 Interfaces) on page 251
show interfaces statistics (for Channelized E1 interfaces) on page 266
show interfaces statistics (for Channelized OC-12 Interfaces) on page 282
show interfaces media (for Channelized STM-1 Interfaces) on page 296
show interfaces statistics (for E1 interfaces) on page 312
show interfaces statistics (for E3 interfaces) on page 323
show interfaces statistics (for ES interfaces) on page 333
show interfaces statistics (for Fast Ethernet interfaces) on page 346
show interfaces statistics (for Gigabit Ethernet interfaces) on page 359
show interfaces statistics (for GRE Interfaces) on page 369
show interfaces statistics (for IP-IP interfaces) on page 379
show interfaces statistics (for Multilink Services Interfaces) on page 391
show interfaces statistics (for SDH/SONET interfaces) on page 408
show interfaces statistics (for T1 Interfaces) on page 419
show interfaces statistics (for T3 interfaces) on page 433
show interfaces statistics (for VT interfaces) on page 443
show interfaces statistics (for ES interfaces) on page 333
show interfaces statistics (for ES interfaces) on page 333
show interfaces statistics (for ES interfaces) on page 333

show interfaces terse

Syntax show interfaces terse <*interface-name*>

Description Display summary information about interfaces.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level view

Output Fields Interface—Name of the interface.

Admin—Whether the interface is turned on (up) or off (down).

Link—Link state. It can be Up or Down.

Proto—Protocol configured on that interface.

Local—Local address of the interface.

Remote—For point-to-point interfaces, address of the remote side of the connection.

Sample Output

```
user@host> show interfaces terse
Interface       Admin Link Proto Local             Remote
at-1/3/0        up    up
at-1/3/0.0      up    up   inet  1.0.0.1           --> 1.0.0.2
                           iso
fxp0            up    up
fxp0.0          up    up   inet  192.168.5.59/24
fxp1            up    up
fxp1.0          up    up   tnp   1
gre             down  up
ipip            down  up
lo0             up    up
lo0.0           up    up   inet  127.0.0.1         --> 0/0
                           iso
47.0005.80ff.f800.0000.0108.0001.1921.6800.5059.00
so-1/2/0        up    down
so-1/2/1        down  down
so-1/2/2        down  down
so-1/2/3        down  down
so-2/0/0        up    up
so-2/0/0.0      up    up   inet  1.2.3.4           --> 1.2.3.5
                           iso
so-2/0/1        down  down
so-2/0/2        down  down
so-2/0/3        down  down
so-2/1/0        down  down
so-2/1/1        down  down
so-2/1/2        down  down
so-2/1/3        down  down
so-2/2/0        down  down
so-2/2/1        down  down
so-2/2/2        down  down
so-2/2/3        down  down
so-2/3/0        down  down
so-4/0/0        down  down
```

```
user@host> show interfaces terse at-1/3/0
Interface       Admin Link Proto Local            Remote
at-1/3/0        up    up
at-1/3/0.0      up    up   inet  1.0.0.1          --> 1.0.0.2
                           iso
user@host> show interfaces ml-1/0/0 terse
Interface       Admin Link Proto Local                Remote
ml-1/0/0        up    up
ml-1/0/0.1      up    up   inet  10.20.30.40       --> 10.20.30.50
ml-1/0/0.2      up    up   inet  10.0.0.2          --> 10.0.0.1
                           iso
                           mpls
ml-1/0/0.10     up    up   inet  128.1.1.3         --> 128.1.1.2
                           iso
```

Chapter 9
ATM Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot ATM interfaces.

show interfaces (for ATM Interfaces)

Syntax

show interfaces at-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces at-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces at-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about ATM router interfaces.

Options

none—Display information about all interfaces.

at-*fpc/pic/port*—Name of an ATM interface.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for ATM Interfaces) on page 217
show interfaces brief (for ATM Interfaces) on page 218
show interfaces detail (for ATM Interfaces) on page 218
show interfaces extensive (for ATM Interfaces) on page 219
show interfaces media (for ATM Interfaces) on page 222
show interfaces statistics (for ATM Interfaces) on page 222

Output Fields at a Glance

Table 17 summarizes which information is included in the output fields of each show interfaces command option for ATM interfaces. In this table, output fields are listed in alphabetical order. In Table 18, the output fields are listed in more detail in the order in which they are displayed.

Table 17: ATM Show Interfaces Output Field Summary in Alphabetical Order

Options	Output Field Description
Physical Interface	
Detail Extensive Standard	Active alarms and Active defects—ATM media-specific defects 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.
Extensive	ATM Encapsulation—ATM encapsulation type.
Extensive	ATM Statistics—ATM statistics for the interface.
Extensive	ATM Status—ATM state information: HCS State, LOC.
All	Clocking—Reference clock source. It can be Internal or External.
Extensive	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.
All	Device Flags Field—Information about the physical device.
Extensive	DS3 media—Counts of T3 media-specific errors.
All	Enabled Field—State of the interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
All	Generation—A unique number for use by Juniper Networks Customer Support only.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	Input errors, Output errors—Detailed input and output errors.
Standard	Input packets, Output packets—Number of packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
Extensive	Line build-out—Configured line buildout setting.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link flags—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	Mode—Whether C-bit parity or M13 mode is enabled.
All	MTU—MTU size on the physical interface.
All	Payload scrambler—Whether payload scrambling is enabled.
Extensive	PFE configuration—Information about the configuration of the Packet Forwarding Engine.
All	Physical interface and Index—Name of the physical interface and the physical interface's index number, which reflects its initialization sequence.
Extensive	PLCP statistics—PLCP error statistics.
Extensive	Queue counters—Name of queue. It can be 0, 1, 2, or 3.
Extensive	Received SONET/SDH overhead, Transmitted SONET/SDH overhead—Values in the received and transmitted SONET/SDH overhead.

Options	Output Field Description
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Extensive	SONET PHY, SONET section, SONET line, SONET path—Counts of specific SONET/SDH errors.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Extensive	Tail queue drops—Number of packets dropped due to bandwidth constraints. Indicates that packets are queued to send out at a rate faster than allowed.
Detail Extensive	Traffic statistics—Number of packets and bytes received and transmitted on the physical interface.
Logical Interface	
All	Addresses—Addresses associated with the logical interface.
All	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Encapsulation—Encapsulation on the logical interface.
All	Family Flags Field—Information about the protocol family flags.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Local—IP address of the logical interface.
All	Logical interface and Index—Name of the logical interface and the logical interface's index number, which reflects its initialization sequence.
All	Logical Interface Flags Field—Information about the logical interface.
All	MTU—MTU size on the logical interface.
Detail Extensive	OAM F5 cell statistics—Number of OAM F5 cells the VCI has sent and received.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
All	VCI—Virtual circuit identifier number and information about the VCI.
Detail Extensive	VCI ARP statistics—Statistics about ARP traffic.
All	VCI Flags—Status of VCI.
Detail Extensive	VCI per-VC transmit statistics—Statistics for ATM VC.
All	VCI Total down time—Number of seconds VCI has been down and time of last Down transition.
All	VCI Traffic statistics—Information about VCI traffic.

Table 18: ATM Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
Payload scrambler	Whether payload scrambling is enabled.
Mode	Whether C-bit parity mode or M13 mode is enabled.
Line build-out	Configured line buildout setting. The default value is 10 feet. distance can be any integer value, followed by a unit specifier of ft or m. The range is 0 through 255 feet for E3 traffic or 0 through 450 feet for T3 traffic.
ATM Encapsulation	ATM encapsulation type. It can be PLCP or Direct. PLCP is the default value.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Current address	Ethernet MAC address for this interface for Ethernet over ATM encapsulation.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
Traffic statistics	Statistics for traffic on the interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable.

Output Field	Output Field Description
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Queue counters	Name of queue. It can be 0, 1, 2, or 3. ■ Queued packets—Number of queued packets. ■ Transmitted packets—Number of transmitted packets. ■ RED-dropped packets—Number of RED-dropped packets.

Active alarms and Active defects	ATM media-specific defects 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. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or the red or yellow alarm LED on the craft interface. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see Chapter 6 in GR-253, Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria.
	■ AIS—Alarm indication signal
	  ■ AIS-L—Alarm indication signal (line)
	  ■ AIS-P—Alarm indication signal (path)
	■ BERR—Bit error rate
	  ■ BERR-SD—Bit error rate defect–signal degrade
	  ■ BERR-SF—Bit error rate fault–signal fail
	■ EXZ—Excessive zeros
	■ FERF—For end receive failures
	■ IDLE— Idle code detected
	■ LCD—Loss of cell delineation
	■ LCV—Line code violation
	■ LOC—Loss of cell delineation
	■ LOF—Loss of frame
	■ LOL—Loss of light
	■ LOP—Loss of pointer
	■ LOS—Loss of signal
	■ PLL—Phase-locked loop out of lock
	■ PLCP_LOF—Loss of PLCP frame alarm
	■ PLCP_YLW PLCP—Alarm at the remote end
	■ PLM-P—Payload label mismatch
	■ RDI—Remote defect indication
	  ■ RDI-L—Remote defect indication (line)
	  ■ RDI-P—Remote defect indication (path)
	■ REI—Remote error indication
	■ SEF—Severely errored frame
	■ UNEQ—Unequipped
	■ YLW—Remote defect indication

Output Field	Output Field Description
■ SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Active alarms and defects, plus counts of specific SONET/SDH errors with detailed information. ■ Seconds—Number of seconds the defect has been active. ■ Count—Number of times that the defect has gone from inactive to active. ■ State—Indicates state of the error. State other than OK indicates a problem. In addition to those listed under Active alarms and Active defect, the SONET PHY, SONET section, SONET line, SONET path fields report the following values: ■ BIP—Bit Interleaved Parity ■ BIP-B1—Bit Interleaved Parity for SONET/SDH section overhead ■ BIP-B2—Bit Interleaved Parity for SONET/SDH line overhead ■ BIP-B3—Bit Interleaved Parity for SONET/SDH path overhead ■ ES—Errored Seconds ■ ES-L—Errored Seconds (near-end line) ■ ES-LFE—Errored Seconds (far-end line) ■ ES-P—Errored Seconds (near-end STS path) ■ ES-PFE—Errored Seconds (far-end STS path) ■ ES-S—Errored Seconds (section) ■ PHY Light—Loss of optical signal ■ REI—Remote Error Indication ■ REI-L—Remote Error Indication (line) ■ REI-P—Remote Error Indication (path) ■ SEF—Severely Errored Frame ■ SEFS-S—Severely Errored Framing Seconds (section) ■ SES—Severely Errored Seconds ■ SES-L—Severely Errored Seconds (near-end line) ■ SES-LFE—Severely Errored Seconds (far-end line) ■ SES-P—Severely Errored Seconds (near-end STS path) ■ SES-PFE—Severely Errored Seconds (far-end STS path) ■ SES-S—Severely Errored Seconds (section) ■ UAS—Unavailable Seconds ■ UAS-L—Unavailable Seconds (near-end line) ■ UAS-LFE—Unavailable Seconds (far-end line) ■ UAS-P—Unavailable Seconds (near-end STS path) ■ UAS-PFE—Unavailable Seconds (far-end STS path)
Received SONET overhead, Transmitted SONET overhead	(Extensive output only) Values in the received and transmitted SONET/SDH overhead.
DS3 media	(Extensive output only) Counts of T3 media-specific errors.
ATM Status	(Extensive output only) ATM state information: ■ HCS State—Shows the status of the header check sequence. ATM uses the HCS field in the cell header in the cell delineation process to frame ATM cell boundaries. The HCS is a FCS-8 calculation over the first four octets of the ATM cell header. ■ LOC—Shows the current loss of cell delineation state. OK means that no LOC is currently asserted.
PLCP defects	■ Display PLCP defects in Seconds, Count, and State.
ATM Status	ATM state information display HCS State and LOC.
PLCP statistics	PLCP error statistics display Framing Errors, Bit interleaved parity errors and far-end block errors.

Output Field	Output Field Description
ATM Statistics	(Extensive output only) ATM statistics for the interface: ■ Uncorrectable HCS errors—Number of cells dropped because the cell delineation failed. These errors most likely indicate that a SONET/SDH layer problem has occurred. ■ Correctable HCS errors—Number of correctable HCS errors that occurred. The cell delineation process can recover from these errors and locate the ATM cell boundary, although the framing process is not quite stable. The ATM cell is not dropped. This counter increases a bit when the cell delineation process changes its state from present to sync (for example, when a fiber is plugged into the interface). The following error statistics are from the framer: ■ Tx cell FIFO overruns—Number of overruns in the transmit FIFO. ■ Rx cell FIFO overruns—Number of overruns in the receive FIFO. ■ Rx cell FIFO underruns—Number of underruns in the receive FIFO. ■ Input cell count—Number of ATM cells received by the interface (not including idle cells). ■ Output cell count—Number of ATM cells transmitted by the interface (not including idle cells). ■ Output idle cell count—Number of idle cells sent by the port. When ATM has nothing to send, it sends idle cells to fill the time slot. ■ Output VC queue drop—Number of packets dropped by a port on the PIC. Packets are dropped because of queue limits on the VCs. The following error statistics are from the SAR: ■ Input no buffers—Number of AAL5 packets dropped because no channel blocks or buffers were available to handle them. ■ Input length errors—Number of AAL5 packets dropped because their length was incorrect. Usually, these errors occur because a cell has been corrupted or lost, or because the length field was corrupted. They can also mean the AAL5 length field was zero. ■ Input timeout—Number of AAL5 packets dropped because of a reassembly timeout. ■ Input invalid vc—Number of AAL5 packets dropped because the header was an unrecognized header (because the VC was not correct or not configured). ■ Input bad CRCs—Number of AAL5 packets dropped because of frame check sequence errors. ■ Input OAM cell no buffer—Number of receive OAM cells or raw cells dropped because no buffers were available to handle them.
Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine: ■ Destination slot—FPC slot number. ■ Stream number—Stream used by the ASIC on the FPC. ■ CoS transmit queue bandwidth, CoS weighted round-robin—Class-of-service parameters. They are not applicable to ATM PICs, and their values should always be 0.
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Protocol	Protocol running on the logical interface.
MTU	MTU size on the logical interface.

Output Field	Output Field Description
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.
VCI	Virtual circuit identifier number and information about the VCI. VBR ■ Peak, Sustained, Burst size, and Queue length. ■ Correlation tag mismatches OAM ■ Period—Configured time interval controlling how often OAM F5 loopback cells are sent. ■ Up count—Number of seconds VCI has been up. ■ Down count—Number of seconds VCI has been down. ■ Total down time—Total number of seconds VCI has been down, and time of last Down transition. ■ Last down— Time of last Down transition. ATM per-VC transmit statistics—Statistics for ATM VC. ■ Tail queue drops—Number of packets dropped due to bandwidth constraints. Indicates that packets are queued to send out at a rate faster than allowed. OAM F5 cell statistics—Number of OAM F5 cells the interface has sent and received. ■ Loop—Number of OAM F5 loopback cells (similar in function to keepalives). ■ RDI—Number of OAM F5 cells with the Remote Defect Indication bit set. ■ AIS—Number of OAM F5 cells with the Alarm Indication Signal bit set. ATM ARP statistics. Statistics about ARP traffic. ■ Received—Number of Inverse ATM ARP messages received. ■ Sent—Number of Inverse ATM ARP messages sent. ■ Denied—Number of times an Inverse ATM ARP response was not sent. ■ Opnotsupp—Number of ATM ARP packets received other than Inverse ARP request messages. ■ Bad len—Bad packet length. ■ Bad proto—Bad protocol (not IP protocol). ■ Bad prlen—Bad protocol length. ■ Bad hrdlen—Bad hardware length.
VCI Multipoint destination	Multipoint IP destination address.
VCI Traffic statistics	Information about VCI traffic.

Output Field	Output Field Description
VCI Flags	Status of VCI. It can be one or more of the following: ■ Active—VCI is up and in working condition. ■ Closed—VCI is closed as a result of user disabling either the logical or physical interface from the CLI. ■ Down—VCI is not in working condition. Some examples are the presence of alarms, defects, F5 AIS/RDI, or no response to F5 PAM loopback cells. ■ ILMI—ILMI is enabled. VCI flags. It can be one or more of the following: ■ Inverse-ARP—ATM Inverse-ARP is enabled. ■ Multicast—VCI is a multicast-VCI (ATM) or multicast DLCI (Frame Relay). ■ OAM—F5 OAM loopback is enabled. For more details, refer to the OAM-period configuration section of the JUNOS Internet Software Configuration Guide: Interfaces and Class of Service. ■ Passive-OAM—Passive-OAM is enabled. When this mode is on, the interface does not respond to F5 OAM AIS/RDI cells received; it only collects statistics. Passive-OAM does not generate OAM loopback request cells even if OAM-period is configured. ■ Shaping—Shaping is enabled.

show interfaces (standard) (for ATM Interfaces)

```
user@host> show interfaces at-6/0/1
Physical interface: at-6/0/1, Enabled, Physical link is Up
  Interface index: 34, SNMP ifIndex: 13
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
  Speed: OC3, Loopback: None, Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None

  Logical interface at-6/0/1.0 (Index 11) (SNMP ifIndex 56)
    Flags: Point-To-Multipoint SNMP-Traps Encapsulation: ATM-SNAP
  Input packets : 0
  Output packets: 0
    VCI 0.128
      Flags: Active, Multicast
      Total down time: 0 sec, Last down: Never
```

show interfaces brief (for ATM Interfaces)

```
user@host> show interfaces at-7/0/0 brief
Physical interface: at-7/0/0, Enabled, Physical link is Up
  Description: singapore51 at-1/1/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC12, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None

  Logical interface at-7/0/0.100
    Flags: Point-To-Point SNMP-Traps Encapsulation: ATM-SNAP
    inet  10.9.140.2       --> 10.9.140.1
    iso
    mpls
    VCI 0.200
      Flags: Active, Inverse-ARP, OAM
      Total down time: 0 sec, Last down: Never
```

show interfaces detail (for ATM Interfaces)

```
user@host> show interfaces at-7/0/0 detail
Physical interface: at-7/0/0, Enabled, Physical link is Up
  Interface index: 101, SNMP ifIndex: 106, Generation: 100
  Description: singapore51 at-1/1/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC12, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:10:c7:72
  Last flapped   : 2002-05-23 12:12:18 PDT (1d 03:12 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              9493424                 1688 bps
   Output bytes  :             10421849                  592 bps
   Input  packets:               129526                    1 pps
   Output packets:               126488                    0 pps
  SONET alarms   : None
  SONET defects  : None

  Logical interface at-7/0/0.100 (Index 49) (SNMP ifIndex 143)
(Generation 76)
    Flags: Point-To-Point Inverse-ARP SNMP-Traps Encapsulation: ATM-SNAP
    Traffic statistics:
     Input  bytes  :                 1039
     Output bytes  :                 3879
     Input  packets:                   17
     Output packets:                   13
    Local statistics:
     Input  bytes  :                 1039
     Output bytes  :                 3879
     Input  packets:                   17
     Output packets:                   13
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
```

```
    Protocol inet, MTU: 4470, Flags: None, Generation: 200 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.9.140.1, Local: 10.9.140.2, Broadcast:
    Unspecified, Generation: 106
    Protocol iso, MTU: 4470, Flags: None, Generation: 201 Route table: 0
    Protocol mpls, MTU: 4458, Flags: None, Generation: 202 Route table: 0
    VCI 0.200
      Flags: Active, Inverse-ARP, OAM, Shaping
      VBR, Peak: 12mbps, Sustained: 10mbps, Burst size: 24, Queue length: 0
      OAM, Period 10 sec, Up count: 5, Down count: 4
      Total down time: 0 sec, Last down: Never
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
     OAM F5 cell statistics:
      Total received: 3, Total sent: 3
      Loopback received: 3, Loopback sent: 3
      Last received: 00:00:07, Last sent: 00:00:07
      RDI received: 0, RDI sent: 0
      AIS received: 0
      Traffic statistics:
       Input  bytes  :                 1039
       Output bytes  :                 3879
       Input  packets:                   17
       Output packets:                   13
```

show interfaces extensive (for ATM Interfaces)

```
user@host> show interfaces at-7/0/0 extensive
Physical interface: at-7/0/0, Enabled, Physical link is Up
  Interface index: 101, SNMP ifIndex: 106, Generation: 100
  Description: singapore51 at-1/1/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC12, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:10:c7:72
  Last flapped   : 2002-05-23 12:12:18 PDT (1d 03:20 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              9526697                  744 bps
   Output bytes  :             10458384                  496 bps
   Input  packets:               129969                    0 pps
   Output packets:               126940                    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
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0
  SONET alarms   : None
  SONET defects  : None
  SONET PHY:            Seconds        Count  State
    PLL Lock                  0            0  OK
    PHY Light                 0            0  OK
```

```
  SONET section:
    BIP-B1                    1            9
    SEF                       0            0  OK
    LOS                       0            0  OK
    LOF                       0            0  OK
    ES-S                      1
    SES-S                     0
    SEFS-S                    0
  SONET line:
    BIP-B2                    1          183
    REI-L                     1          323
    RDI-L                     0            0  OK
    AIS-L                     0            0  OK
    BERR-SF                   0            0  OK
    BERR-SD                   0            0  OK
    ES-L                      1
    SES-L                     0
    UAS-L                     0
    ES-LFE                    1
    SES-LFE                   0
    UAS-LFE                   0
  SONET path:
    BIP-B3                    1           31
    REI-P                     1          216
    LOP-P                     0            0  OK
    AIS-P                     0            0  OK
    RDI-P                     0            0  OK
    UNEQ-P                    0            0  OK
    PLM-P                     0            0  OK
    ES-P                      1
    SES-P                     0
    UAS-P                     0
    ES-PFE                    1
    SES-PFE                   0
    UAS-PFE                   0
  Received SONET overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x13, C2(cmp) : 0x13, F2      : 0x00
    Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x00
    V5(cmp) : 0x00
  Transmitted SONET overhead:
    F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x13, F2      : 0x00, Z3      : 0x00
    Z4      : 0x00, V5      : 0x00
  ATM status:
    HCS state:     Sync
    LOC      :       OK
  ATM Statistics:
    Uncorrectable HCS errors: 77, Correctable HCS errors: 5, Tx cell
FIFO overruns: 0,
    Rx cell FIFO overruns: 1, Rx cell FIFO underruns: 0, Input cell
count: 421980,
    Output cell count: 139110927341, Output idle cell count: 1671702365,
Output VC queue drops: 0,
    Input no buffers: 0, Input length errors: 0, Input timeouts: 0,
Input invalid VCs: 143301,
    Input bad CRCs: 0, Input OAM cell no buffers: 0
  Packet Forwarding Engine configuration:
    Destination slot: 7
    CoS transmit queue          Bandwidth          Buffer     Priority
```

```
Limit
                              %          bps   %        bytes
    0 best-effort             0            0   0            0
low    none
    1 expedited-forwarding    0            0   0            0
low    none
    2 assured-forwarding      0            0   0            0
low    none
    3 network-control         0            0   0            0
low    none
Logical interface at-7/0/0.100 (Index 49) (SNMP ifIndex 143) (Generation 76)
    Flags: Point-To-Point Inverse-ARP SNMP-Traps Encapsulation: ATM-SNAP
    Traffic statistics:
     Input  bytes  :                 9993
     Output bytes  :                16246
     Input  packets:                  151
     Output packets:                  136
    Local statistics:
     Input  bytes  :                 9993
     Output bytes  :                16246
     Input  packets:                  151
     Output packets:                  136
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4470, Flags: None, Generation: 200 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.9.140.1, Local: 10.9.140.2, Broadcast:
    Unspecified, Generation: 106
    Protocol iso, MTU: 4470, Flags: None, Generation: 201 Route table: 0
    Protocol mpls, MTU: 4458, Flags: None, Generation: 202 Route table: 0
    VCI 0.200
      Flags: Active, Inverse-ARP, OAM, Shaping
      VBR, Peak: 12mbps, Sustained: 10mbps, Burst size: 24, Queue length: 0
      OAM, Period 10 sec, Up count: 5, Down count: 4
      Total down time: 0 sec, Last down: Never
      ATM per-VC transmit statistics:
       Tail queue packet drops: 0
     OAM F5 cell statistics:
      Total received: 49, Total sent: 49
      Loopback received: 49, Loopback sent: 49
      Last received: 00:00:08, Last sent: 00:00:08
      RDI received: 0, RDI sent: 0
      AIS received: 0
      Traffic statistics:
       Input  bytes  :                 9993
       Output bytes  :                16246
       Input  packets:                  151
       Output packets:                  136
```

show interfaces media (for ATM Interfaces)

```
user@host> show interfaces media at-7/0/0
Physical interface: at-7/0/0, Enabled, Physical link is Up
  Interface index: 101, SNMP ifIndex: 106
  Description: singapore51 at-1/1/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC12, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Current address: 00:90:69:10:c7:72
  Last flapped   : 2002-05-23 12:12:18 PDT (1d 03:22 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 296 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
    BIP-B1: 9, BIP-B2: 183, REI-L: 323, BIP-B3: 31, REI-P: 216
    BIP-BIP2: 0
```

show interfaces statistics (for ATM Interfaces)

```
user@host> show interfaces statistics at-7/0/0
Physical interface: at-7/0/0, Enabled, Physical link is Up
  Interface index: 101, SNMP ifIndex: 106
  Description: singapore51 at-1/1/0
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, SONET mode,
Speed: OC12, Loopback: None,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Link flags     : None
  Current address: 00:90:69:10:c7:72
  Last flapped   : 2002-05-23 12:12:18 PDT (1d 03:23 ago)
  Statistics last cleared: Never
  Input rate     : 592 bps (0 pps)
  Output rate    : 296 bps (0 pps)
  Input errors: 0, Output errors: 0
  SONET alarms   : None
  SONET defects  : None

  Logical interface at-7/0/0.100 (Index 49) (SNMP ifIndex 143)
    Flags: Point-To-Point Inverse-ARP SNMP-Traps Encapsulation: ATM-SNAP
  Input packets : 196
  Output packets: 177
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.9.140.1, Local: 10.9.140.2
    Protocol iso, MTU: 4470, Flags: None
    Protocol mpls, MTU: 4458, Flags: None
    VCI 0.200
      Flags: Active, Inverse-ARP, OAM, Shaping
      VBR, Peak: 12mbps, Sustained: 10mbps, Burst size: 24, Queue length: 0
      OAM, Period 10 sec, Up count: 5, Down count: 4
      Total down time: 0 sec, Last down: Never
     OAM F5 cell statistics:
      Total received: 65, Total sent: 65
      Loopback received: 65, Loopback sent: 65
      RDI received: 0, RDI sent: 0
      AIS received: 0
```

```
Traffic statistics:
     Input  packets:                  196
     Output packets:                  177
```

Chapter 10

Channelized DS-3 to DS-0 Interfaces Monitoring and Troubleshooting

This chapter displays the command-line interface (CLI) commmands you can use to monitor and troubleshoot Channelized DS-3 to DS-0 router interfaces.

show interfaces (for Channelized DS-3 to DS-0 Interfaces)

Syntax

show interfaces ds-*fpc/pic/port:T1channel:DS-0 channel* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces ds-*fpc/pic/port:T1channel:DS-0 channel* <brief | detail | extensive> <media>

show interfaces ds-*fpc/pic/port:T1channel:DS-0 channel* <brief | detail | extensive> <statistics>

Description

Display status information about Channelized DS-3 to DS-0 router interfaces.

Options

none—Display information about all interfaces.

ds-*fpc/pic/port:T1channel:DS-0 channel* —Name of a DS-3 to DS-0 interface.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Channelized DS-3 to DS-0 Interfaces) on page 233
show interfaces brief (for Channelized DS-3 to DS-0 Interfaces) on page 233
show interfaces detail (for Channelized DS-3 to DS-0 Interfaces) on page 234
show interfaces extensive (for Channelized DS-3 to DS-0 Interfaces) on page 235
show interfaces media (for Channelized DS-3 to DS-0 Interfaces) on page 237
show interfaces statistics (for Channelized DS-3 to DS-0 Interfaces) on page 237

Output Fields at a Glance

Table 19 summarizes which information is included in the output fields of each show interfaces command option for Channelized DS-3 to DS-0 interfaces. In this table, output fields are listed in alphabetical order. In Table 20, the output fields are listed in the order in which they are displayed.

Table 19: Channelized DS-3 to DS-0 Show Interfaces Output Field Summary in Alphabetical Order

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—T1 media-specific defects that can render the interface unable to pass packets.
All	Active alarms and Active defects—T3 media-specific defects that can render the interface unable to passpackets.
Detail Extensive	BERT time period—Configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
All	Device Flags Field—Information about the physical device.
Extensive	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.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—BERT (Bit Error Rate Testing) checks the quality of the line.
Extensive	DS3 media—Counts of T3 media-specific errors.
Extensive	DSU configuration—Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
All	Enabled Field—State of the interface.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16 bit.
All	Framing—Physical layer framing format used on the link. It can be ESF or SF. The default is ESF.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.

Options	Field Description
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	Interface transmit queues—Transmit queue statistics and configuration for each DS-0 channel on the Channelized DS-3 to DS-0 PIC.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Detail Extensive	Keepalive settings—Configure settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
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).
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	Mode—Whether C-bit parity mode is enabled.
All	MTU—MTU size on the physical interface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Extensive	T1 media—Counts of T1 media-specific errors.
Extensive	T3 media—Counts of T3 media-specific errors.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Logical Interface	
All	Addresses flags—Information about the address.
All	Addresses—Addresses associated with the logical interface.
Detail Extensive	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
All	DLCI—If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
All	Encapsulation—Encapsulation on the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.

Options	Field Description
All	Flags—Information about the logical interface.
Standard	Input packets, Output packets—Number of packets received and transmitted on the logical interface.
Detail Extensive	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine.
All	Local—IP address of the logical interface.
All	Logical interface flags—Information about the logical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number, and the logical interface's SNMP interface index number.
Detail Extensive	MTU Family Flags Field—Information about the protocol family flags.
All	MTU—MTU size on the logical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics.

Table 20: Channelized DS-3 to DS-0 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
Mode	Whether C-bit parity mode or M13 mode is enabled.
FCS	Frame check sequence on the interface (either 16 or *32*). The default is 16 bit.
Framing	Physical layer framing format used on the link. It can be ESF or SF. The default is ESF.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface. Possible values are described in "Interface Flags Field" on page 191.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)

Output Field	Output Field Description
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Keepalive statistics	Information about keepalive packets.
	■ Input—Number of keepalive packets received by PPP.
	  ■ (last seen 00:00:00 ago)—Time the last keepalive packet was received in the format *hh:mm:ss*.
	■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ■ (last seen 00:00:00 ago)—Time the last keepalive packet was sent in the format *hh:mm:ss*.

Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with Radius authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. Input bytes, Output bytes Number of bytes received and transmitted on the interface. Input packets, Output packets Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits and packets received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects (T3)	T3 media-specific defects 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. 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. ■ AIS—Alarm indication signal ■ EXZ—Excessive zeros ■ FERF—Far-end receive failures ■ IDLE— Idle code detected ■ LCV—Line code violation ■ LOS—Loss of signal ■ LOF—Loss of frame ■ PLL—Phase locked loop ■ YLW—Remote defect indication ■ T3 media—Counts of T3 media-specific errors

Output Field	Output Field Description
Active alarms and Active defects (T1)	T1 media-specific defects 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. 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. ■ LOS—Loss of signal. ■ LOF—Loss of frame. ■ AIS—Alarm indication signal. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.
T1 media	Counts of T1 media-specific errors.
DS3 media	(Extensive output only) Counts of T3 media-specific errors.
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It can be B8ZS or AMI. ■ Byte encoding—Byte encoding used. It can be Nx64K or Nx56K. ■ Data inversion—HDLC data inversion setting. It can be Enabled or Disabled.
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
DS3 BERT, DS1 BERT configuration	BERT (Bit Error Rate Test) checks the quality of the line. This output only appears when BERT is run on the interface (see “test interface bert-start” on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT. It can be Pseudorandom, Repetitive, or All ones.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only non-zero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Input packets, Output packets	(Standard output only) Number of packets received and transmitted on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Protocol	Protocol running on the logical interface.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.
Interface transmit queues	Transmit queue statistics and configuration for each DS-0 channel on the Channelized DS-3 to DS-0 PIC.

show interfaces (standard) (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0
 Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 278 (00:00:05 ago), Output: 283 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:46:23 ago)
  Input rate     : 16 bps (0 pps)
  Output rate    : 72 bps (0 pps)
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None

  Logical interface ds-0/0/0:0:0.0 (Index 5) (SNMP ifIndex 4299)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 18.18.18.1, Local: 18.18.18.2
```

show interfaces brief (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0 brief
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 279 (00:00:03 ago), Output: 284 (00:00:05 ago)

  Logical interface ds-0/0/0:0:0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  18.18.18.2       --> 18.18.18.1
```

show interfaces detail (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0 detail
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298, Generation: 177
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 280 (last seen 00:00:02 ago)
    Output: 285 (last sent 00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:46:39 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 6790                   16 bps
   Output bytes  :                28748                   72 bps
   Input  packets:                  566                    0 pps
   Output packets:                  890                    0 pps
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  DS-3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: Unknown (0), Induced error rate: 10e-0
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)

  Logical interface ds-0/0/0:0:0.0 (Index 5) (SNMP ifIndex 4299)
  (Generation 943)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 949, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 18.18.18.1, Local: 18.18.18.2, Broadcast: Unspecified,
        Generation: 1849
```

show interfaces extensive (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0 extensive
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298, Generation: 177
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 280 (last seen 00:00:09 ago)
    Output: 286 (last sent 00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:46:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 6814                   16 bps
   Output bytes  :                28840                   72 bps
   Input  packets:                  568                    0 pps
   Output packets:                  893                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 39, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 2, L2 mismatch timeouts: 0,
    HS link CRC errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  T1  media:            Seconds        Count  State
    SEF                       0            0  OK
    BEE                       5            1  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
    YELLOW                   17            1  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       5        27765
    PCV                       0            0
    CS                        0            0
    LES                       0
    ES                        0
    SES                       5
    SEFS                     10
    BES                       0
    UAS                       0
```

```
  DS3 media:            Seconds        Count  State
    PLL Lock                  0            0  OK
    Reframing                 0            0  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
    IDLE                      0            0  OK
    YELLOW                    0            0  OK
    BPV                       1        65535
    EXZ                       1        65535
    LCV                       2       131070
    PCV                       1         1825
    CCV                       0            0
    LES                       1
    PES                       1
    PSES                      1
    CES                       0
    CSES                      0
    SEFS                      0
    UAS                       0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5           893        28840            0            0
  HDLC configuration:
    Giant threshold: 1514, Runt threshold: 3
    Timeslots      : 1-10
    Byte encoding: Nx64K, Data inversion: Disabled
  DS-3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: Unknown (0), Induced error rate: 10e-0
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 2 (0x01)
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                              %          bps   %        bytes
    0 best-effort            95       608000  95            0      low    none
    3 network-control         5        32000   5            0      low    none

 Logical interface ds-0/0/0:0:0.0 (Index 5) (SNMP ifIndex 4299)
  (Generation 943)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 949, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 18.18.18.1, Local: 18.18.18.2, Broadcast: Unspecified,
       Generation: 1849
```

show interfaces media (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0 media
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 281 (00:00:10 ago), Output: 287 (00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:46:58 ago)
  Input rate     : 16 bps (0 pps)
  Output rate    : 72 bps (0 pps)
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  T1 errors:
    BPV: 0, EXZ: 0, LCV: 27765, PCV: 0, CS: 0
  DS3 errors:
    BPV: 65535, EXZ: 65535, LCV: 131070, PCV: 1825, CCV: 0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5           896        28932            0            0
```

show interfaces statistics (for Channelized DS-3 to DS-0 Interfaces)

```
user@host> show interfaces ds-0/0/0:0:0 statistics
Physical interface: ds-0/0/0:0:0, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 4298
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps, FCS: 16,
  Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 282 (00:00:07 ago), Output: 287 (00:00:12 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:53:29 PDT (00:47:07 ago)
  Statistics last cleared: Never
  Input rate     : 16 bps (0 pps)
  Output rate    : 72 bps (0 pps)
  Input errors: 41, Output errors: 0
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
```

```
Logical interface ds-0/0/0:0:0.0 (Index 5) (SNMP ifIndex 4299)
  Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 18.18.18.1, Local: 18.18.18.2
```

Chapter 11

Channelized DS-3 to DS-1 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Channelized DS-3 to DS-1 interfaces.

show interfaces (for Channelized DS-3 to DS-1 Interfaces)

Syntax

show interfaces t1-*fpc/pic/port:channel* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces t1-*fpc/pic/port:channel* <brief | detail | extensive> <media>

show interfaces t1-*fpc/pic/port:channel* <brief | detail | extensive> <statistics>

Description

Display status information about Channelized DS-3 to DS-1 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

t1-*fpc/pic/port:channel*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) for Channelized DS-3 to DS-1 Interfaces) on page 246
show interfaces brief for Channelized DS-3 to DS-1 Interfaces) on page 246
show interfaces detail for Channelized DS-3 to DS-1 Interfaces) on page 247
show interfaces extensive for Channelized DS-3 to DS-1 Interfaces) on page 248
show interfaces media for Channelized DS-3 to DS-1 Interfaces) on page 250
show interfaces statistics for Channelized DS-3 to DS-1 Interfaces) on page 251

Output Fields at a Glance

Table 21 summarizes which information is included in the output fields of each show interfaces command option for Channelized DS-3 to DS-1 interfaces. In this table, output fields are listed in alphabetical order. In Table 22, the output fields are listed in the order in which they are displayed.

Table 21: Channelized DS-3 to DS-1 Show Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—T1 media-specific defects that can render the interface unable to pass packets.
Detail Extensive	BERT time period—Configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—BERT (Bit Error Rate Testing) checks the quality of the line.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16 bit.
All	Device Flags Field—Information about the physical device.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	Framing—Physical layer framing format used on the link. It can be ESF or SF. The default is ESF.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—Instance number of the DS1 interface.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	Interface transmit queues—Transmit queue statistics and configuration for each DS-1 channel on the Channelized DS-3 to DS-1 PIC.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.

Options	Field Description
Detail Extensive	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Extensive	T1 media—Counts of T1 media-specific errors.
Extensive	T3 media—Counts of T3 media-specific errors.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 22: Channelized DS-3 to DS-1 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the Link Management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Keepalive statistics	Information about keepalive packets.
	■ Input—Number of keepalive packets received by PPP.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*.
	■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.

Output Field	Output Field Description
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with Radius authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—Counter increments when the software could not find a valid logical interface (such as t1-1/2/3.0) for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that causes the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces.

Output Field	Output Field Description
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects (T1)	T1 media-specific defects 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. 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. LOS—Loss of signal. LOF—Loss of frame. AIS—Alarm indication signal. YLW—Yellow alarm. Indicates errors at the remote site receiver.
T1 media	Counts of T1 media-specific errors.
T3 media	Counts of T3 media-specific errors.
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It can be B8ZS or AMI. ■ Byte encoding—Byte encoding used. It can be Nx64K or Nx56K. ■ Data inversion—HDLC data inversion setting. It can be Enabled or Disabled. ■ Buildout—Buildout setting.
Interface transmit queues	Name of the transmit queues and their associated statistics for each DS-1 channel on the Channelized DS-3 to DS-1 PIC. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.
DSU Configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
Buildout	Buildout setting.
DS3 BERT, DS1 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see "test interface bert-start" on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT. It can be Pseudorandom, Repetitive, or All ones.

Output Field	Output Field Description

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only non-zero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 26 (00:00:07 ago), Output: 24 (00:00:12 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:30:12 PDT (17:29:01 ago)
  Input rate     : 16 bps (0 pps)
  Output rate    : 16 bps (0 pps)
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None

  Logical interface t1-0/0/0:0.0 (Index 11) (SNMP ifIndex 23)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.21.21.2, Local: 21.21.21.1
```

show interfaces brief for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0 brief
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 27 (00:00:11 ago), Output: 26 (00:00:05 ago)

  Logical interface t1-0/0/0:0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  21.21.21.1       --> 21.21.21.2
```

show interfaces detail for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0 detail
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14, Generation: 2977
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 28 (last seen 00:00:08 ago)
    Output: 27 (last sent 00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:30:12 PDT (17:29:26 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  896                   16 bps
   Output bytes  :                 1114                   16 bps
   Input  packets:                   62                    0 pps
   Output packets:                   78                    0 pps
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  DS-3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: Unknown (0), Induced error rate: 10e-0
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)

  Logical interface t1-0/0/0:0.0 (Index 11) (SNMP ifIndex 23) (Generation 497)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500, Generation: 576, Route table: 0
      Flags: None
        Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.21.21.2, Local: 21.21.21.1, Broadcast: Unspecified,
        Generation: 977
```

show interfaces extensive for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0 extensive
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14, Generation: 2977
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 30 (last seen 00:00:05 ago)
    Output: 29 (last sent 00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:30:12 PDT (17:29:43 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  944                   16 bps
   Output bytes  :                 1162                   16 bps
   Input  packets:                   66                    0 pps
   Output packets:                   82                    0 pps
  Input errors:
    Errors: 1, Drops: 0, Framing errors: 1, Policed discards: 8,
    L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0,
    HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  T1  media:            Seconds        Count  State
    SEF                       0            0  OK
    BEE                      11            5  OK
    AIS                      28            1  OK
    LOF                      27            1  OK
    LOS                       0            0  OK
    YELLOW                   23            1  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                      11        20574
    PCV                       0            0
    CS                        0            0
    LES                      28
    ES                       28
    SES                      39
    SEFS                     50
    BES                       0
    UAS                       0
```

```
DS3 media:            Seconds        Count  State
  PLL Lock                  0            0  OK
  Reframing                 0            0  OK
  AIS                       0            0  OK
  LOF                       1            1  OK
  LOS                       1            1  OK
  IDLE                      0            0  OK
  YELLOW                    0            0  OK
  BPV                       2       131070
  EXZ                       3        49910
  LCV                       5       180980
  PCV                       2          327
  CCV                      12       264558
  LES                       3
  PES                       3
  PSES                      2
  CES                      13
  CSES                     13
  SEFS                      1
  UAS                      35
Interface transmit queues:
            B/W  WRR       Packets        Bytes        Drops       Errors
  Queue0     95   95             0            0            0            0
  Queue1      5    5            82         1162            0            0
HDLC configuration:
  Giant threshold: 1514, Runt threshold: 3
  Timeslots      : 1-10
  Line encoding: B8ZS, Byte encoding: Nx64K, Data inversion: Disabled
DS-3 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Algorithm: Unknown (0), Induced error rate: 10e-0
DS1 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 2 (0x00)
  CoS transmit queue             Bandwidth           Buffer Priority   Limit
                           %          bps   %        bytes
  0 best-effort           95       608000  95            0      low    none
  3 network-control        5        32000   5            0      low    none

Logical interface t1-0/0/0:0.0 (Index 11) (SNMP ifIndex 23) (Generation 497)
  Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
  Bandwidth: 0
  Protocol inet, MTU: 1500, Generation: 576, Route table: 0
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 21.21.21.2, Local: 21.21.21.1, Broadcast: Unspecified,
      Generation: 977
```

show interfaces media for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0 media
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 31 (00:00:03 ago), Output: 30 (00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:30:12 PDT (17:29:52 ago)
  Input rate     : 24 bps (0 pps)
  Output rate    : 24 bps (0 pps)
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None
  T1 errors:
    BPV: 0, EXZ: 0, LCV: 20574, PCV: 0, CS: 0
  DS3 errors:
    BPV: 131070, EXZ: 49910, LCV: 180980, PCV: 327, CCV: 264558
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5            85         1198            0            0
```

show interfaces statistics for Channelized DS-3 to DS-1 Interfaces)

```
user@host> show interfaces t1-0/0/0:0 statistics
Physical interface: t1-0/0/0:0, Enabled, Physical link is Up
  Interface index: 210, SNMP ifIndex: 14
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 640kbps,
  Loopback: None, FCS: 16, Mode: C/Bit parity, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 31 (00:00:09 ago), Output: 31 (00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:30:12 PDT (17:29:58 ago)
  Statistics last cleared: Never
  Input rate     : 24 bps (0 pps)
  Output rate    : 24 bps (0 pps)
  Input errors: 3, Output errors: 0
  DS1   alarms   : None
  DS3   alarms   : None
  DS1   defects  : None
  DS3   defects  : None

  Logical interface t1-0/0/0:0.0 (Index 11) (SNMP ifIndex 23)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 21.21.21.2, Local: 21.21.21.1
```

Chapter 12

Channelized E-1 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Channelized E1 interfaces.

show interfaces (for Channelized E1 Interfaces)

Syntax

show interfaces ds-*fpc/pic/port:ds-0 channel* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces ds-*fpc/pic/port:ds-0 channel* <brief | detail | extensive> <media>

show interfaces ds-*fpc/pic/port:ds-0 channel* <brief | detail | extensive> <statistics>

Description

Display status information about Channelized E1 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

ds-*fpc/pic/port:ds0 channel*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Channelized E1 interfaces) on page 261
show interfaces brief (for Channelized E1 interfaces) on page 262
show interfaces detail (for Channelized E1 interfaces) on page 262
show interfaces extensive (for Channelized E1 interfaces) on page 264
show interfaces media (for Channelized E1 interfaces) on page 266
show interfaces statistics (for Channelized E1 interfaces) on page 266

Output Fields at a Glance

Table 23 summarizes which information is included in the output fields of each show interfaces command option for Channelized E1 interfaces. In this table, output fields are listed in alphabetical order. In Table 24, the output fields are listed in more detail in the order in which they are displayed.

Table 23: Channelized E1 Show Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value, value... xx* seconds
All	Active alarms and Active defects—E1 media-specific defects that can render the interface unable to pass packets.
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
All	Buildout—Buildout setting.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—BERT (Bit Error Rate Testing) checks the quality of the line.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Extensive	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.
All	Device Flags Field—Information about the physical device.
Extensive	E1 media—Counts of E1 media-specific errors.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16 bit.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	Framing—Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—Instance number of the DS0 interface.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value, value... xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.

Options	Field Description
Extensive	Interface transmit queues—Transmit queue statistics and configuration for each DS-0 channel on the Channelized E1 to DS-0 PIC.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface. The default MTU size is 1504.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 24: Channelized E1 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Output Field	Output Field Description
Keepalive statistics	Information about keepalive packets. ■ Input—Number of keepalive packets received by PPP. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*. ■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects	E1 media-specific defects 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. 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. ■ LOS—Loss of signal. ■ LOF—Loss of frame. ■ AIS—Alarm indication signal. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.

Output Field	Output Field Description
E1 media	(Extensive output only) Active alarms and defects, plus counts of specific E1 errors with detailed information. ■ Seconds—Number of seconds the defect has been active. ■ Count—Number of times that the defect has gone from inactive to active. ■ State—State of the error. State other than OK indicates a problem. The E1 media specific error types can be: ■ AIS—Alarm indication signal ■ BEE—Bit error ■ BES—Bit error seconds ■ BPV—Bipolar violation ■ CS—Carrier state ■ ES—Errored Seconds ■ EXZ—Excessive zeros ■ FEBE—far-end block error ■ LCV—Line code violation ■ LES—Line error seconds ■ LOF—Loss of frame ■ LOS—Loss of signal ■ PCV—Pulse code violation ■ SEF—Severely Errored Frame ■ SEFS-S—Severely Errored Framing Seconds (section) ■ SES—Severely Errored Seconds ■ UAS—Unavailable Seconds ■ YELLOW—Errors at the remote site receiver
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Buildout	Buildout setting.
Interface transmit queues	Names of the transmit queues and their associated statistics for each DS0 channel on the Channelized E1 to DS0 PIC. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.
DS1 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see “test interface bert-start” on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT.

Output Field	Output Field Description

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For exampe, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.

Output Field	Output Field Description
Flags	Information about the address flags. Possible values are described in “Address Flags Field” on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 944
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 12 (00:00:10 ago), Output: 13 (00:00:05 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 15:29:39 PDT (00:10:35 ago)
  Input rate     : 293240 bps (32 pps)
  Output rate    : 294456 bps (35 pps)
  DS1   alarms   : None
  DS1   defects  : None

  Logical interface ds-4/3/0:0.0 (Index 200) (SNMP ifIndex 945)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc
    Protocol mpls, MTU: 1488
      Flags: Is-Primary
```

show interfaces brief (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0 brief
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 13 (00:00:02 ago), Output: 13 (00:00:08 ago)

  Logical interface ds-4/3/0:0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  10.41.251.1/24
    iso
    inet6 2060::251:1/126
          fe80::2a0:a5ff:fe12:7dc/64
    mpls
```

show interfaces detail (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0 detail
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 944, Generation: 1429
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 13 (last seen 00:00:05 ago)
    Output: 14 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 15:29:39 PDT (00:10:41 ago)
  Statistics last cleared: 2002-04-26 15:38:07 PDT (00:02:13 ago)
  Traffic statistics:
   Input  bytes  :              4662574               285024 bps
   Output bytes  :              4691918               290584 bps
   Input  packets:                 4567                   32 pps
   Output packets:                 4570                   35 pps
  DS1   alarms   : None
  DS1   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^11 - 1, O.152 and O.153 (2047
```

```
type),Pseudorandom (8)
Logical interface ds-4/3/0:0.0 (Index 200) (SNMP ifIndex 945)
  (Generation 1239)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
    Generation: 3706 Route table: 0
      Filters: Input: sample-all
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1, Broadcast: Unspecified,
        Generation: 6166
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500Generation: 3707 Route table: 0
      Flags: None
    Generation: 3708 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1,
        Broadcast: Unspecified, Generation: 6168
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc,
        Broadcast: Unspecified, Generation: 6169
    Protocol mpls, MTU: 1488
      Flags: Is-Primary
    Generation: 3709 Route table: 0
```

show interfaces extensive (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0 extensive
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 944, Generation: 1429
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 13 (last seen 00:00:10 ago)
    Output: 14 (last sent 00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 15:29:39 PDT (00:10:46 ago)
  Statistics last cleared: 2002-04-26 15:38:07 PDT (00:02:18 ago)
  Traffic statistics:
   Input  bytes  :              4831225               281504 bps
   Output bytes  :              4880158               290584 bps
   Input  packets:                 4733                   32 pps
   Output packets:                 4753                   35 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS link CRC errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  DS1   defects  : None
  E1  media:            Seconds        Count  State
    SEF                       0            0  OK
    BEE                       0            0  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CS                        0            0
    FEBE                      0            0
    LES                       0
    ES                        0
    SES                       0
    SEFS                      0
    BES                       0
    UAS                       0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95            0            0            0            0
    Queue1      5    5         4838      4958934            0            0
  HDLC configuration:
    Giant threshold: 1514, Runt threshold: 3
    Timeslots      : 1-15
    Line encoding: HDB3
```

```
DS1 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^11 - 1, O.152 and O.153 (2047 type),
  Pseudorandom (8)
Packet Forwarding Engine configuration:
  Destination slot: 4, PLP byte: 2 (0x01)
  CoS transmit queue             Bandwidth           Buffer Priority   Limit
                            %          bps   %        bytes
  0 best-effort            95       912000  95            0      low    none
  3 network-control         5        48000   5            0      low    none

Logical interface ds-4/3/0:0.0 (Index 200) (SNMP ifIndex 945)
(Generation 1239)
  Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
  Protocol inet, MTU: 1500
    Flags: None
  Generation: 3706 Route table: 0
    Filters: Input: sample-all
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 10.41.251/24, Local: 10.41.251.1, Broadcast: Unspecified,
      Generation: 6166
  Protocol iso, MTU: 1500
    Flags: None
  Protocol inet6, MTU: 1500Generation: 3707 Route table: 0
    Flags: None
  Generation: 3708 Route table: 0
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 2060::251:0/126, Local: 2060::251:1,
      Broadcast: Unspecified, Generation: 6168
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc,
      Broadcast: Unspecified, Generation: 6169
  Protocol mpls, MTU: 1488
    Flags: Is-Primary
  Generation: 3709 Route table: 0
```

show interfaces media (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0 media
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 944
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 14 (00:00:05 ago), Output: 15 (00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 15:29:39 PDT (00:10:52 ago)
  Input rate     : 281672 bps (32 pps)
  Output rate    : 286760 bps (35 pps)
  DS1   alarms   : None
  DS1   defects  : None
  E1 errors:
    BPV: 0, EXZ: 0, LCV: 0, PCV: 0, CS: 0, FEBE: 0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5          5182      5309300            0            0
```

show interfaces statistics (for Channelized E1 interfaces)

```
user@host> show interfaces ds-4/3/0:0 statistics
Physical interface: ds-4/3/0:0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 944
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: 960kbps, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 14 (00:00:08 ago), Output: 15 (00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 15:29:39 PDT (00:10:55 ago)
  Statistics last cleared: 2002-04-26 15:38:07 PDT (00:02:27 ago)
  Input rate     : 285192 bps (32 pps)
  Output rate    : 294456 bps (35 pps)
  Input errors: 0, Output errors: 0
  DS1   alarms   : None
  DS1   defects  : None

  Logical interface ds-4/3/0:0.0 (Index 200) (SNMP ifIndex 945)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1
    Protocol iso, MTU: 1500
      Flags: None
```

```
Protocol inet6, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 2060::251:0/126, Local: 2060::251:1
  Addresses, Flags: Is-Preferred
    Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc
Protocol mpls, MTU: 1488
  Flags: Is-Primary
```

Chapter 13 Channelized OC-12 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Channelized OC-12 interfaces.

show interfaces (for Channelized OC-12 Interfaces)

Syntax

show interfaces t3-*fpc/pic/port:channel* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces t3-*fpc/pic/port:channel* <brief | detail | extensive> <media>

show interfaces t3-*fpc/pic/port:channel* <brief | detail | extensive> <statistics>

Description

Display status information about Channelized OC-12 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

t3-*fpc/pic/port:channel*—Name of an interface, including channel number.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Channelized OC-12 Interfaces) on page 277
show interfaces brief (for Channelized OC-12 Interfaces) on page 278
show interfaces detail (for Channelized OC-12 Interfaces) on page 278
show interfaces extensive (for Channelized OC-12 Interfaces) on page 279
show interfaces media (for Channelized OC-12 Interfaces) on page 281
show interfaces statistics (for Channelized OC-12 Interfaces) on page 282

Output Fields at a Glance

Table 25 summarizes which information is included in the output fields of each show interfaces command option for Channelized OC-12 interfaces. In this table, output fields are listed in alphabetical order. In Table 26, the output fields are listed in the order in which they are displayed.

Table 25: Channelized OC-12 Show Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value, value... xx* seconds
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—SONET/SDH reference clock source. It can be internal or external. Clocking is configured and displayed only for channel 0.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Test.
All	DS3 alarms and defects—T3 media-specific defects 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. 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.
Detail Extensive	DS3 BERT configuration—The BERT (Bit Error Rate Test) checks the quality of the line.
All	DS3 Loopback—Whether DS-3 loopback is enabled and the type of loopback (local or remote).
Extensive	DS3 media—Counts of T3 media-specific errors.
Extensive	DSU configuration—Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Device flags—Information about the physical device. Possible values are described in “Device Flags Field” on page 190.
Extensive	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.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled—State of the interface. Possible values are described in “Enabled Field” on page 190.
All	FCS—Frame check sequence on the interface (either 16 or 32).
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
All	Framing mode—Framing mode. It can be SONET mode or SDH mode.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value, value... xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.

Options	Field Description
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface flags—Information about the interface. Possible values are described in "Interface Flags Field" on page 191.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	Interface transmit queues—Transmit queue statistics and configuration for each DS-3 channel on the Channelized OC-12 PIC.
Standard Detail	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	Keepalive settings—Configured settings for keepalive packets.
Standard Detail	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
All	Link flags—Information about the link. Possible values are described in "Link Flags Field" on page 191.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (local or remote).
All	Mode—Whether C-bit parity or M13 mode is enabled.
All	MTU—MTU size on the physical interface.
Standard Detail	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Standard Detail	NCP state, LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
Extensive	Received path trace, Transmitted path trace—Channelized OC-12 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. The received path trace value is the message received from the router at the other end of the fiber, and the transmitted path trace value is the message that this router transmits. This information is specific to each of the twelve Channelized OC-12 interfaces.
Extensive	Received SONET overhead, Transmitted SONET overhead—Values in the received and transmitted SONET/SDH overhead.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
All	SONET alarms and SONET defects—POS media-specific defects 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.
All	SONET Loopback—Whether SONET/SDH loopback is enabled and the type of loopback (local or remote).
Extensive	SONET PHY, SONET section, SONET line, SONET path—Counts of specific SONET/SDH errors.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 26: Channelized OC-12 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	SONET/SDH reference clock source. It can be internal or external. Clocking is configured and displayed only for channel 0.
Framing mode	Framing mode can be SONET or SDH.
Speed	Speed at which the interface is running.
DS3 Loopback	Whether DS-3 loopback is enabled and the type of loopback (local or remote).
SONET Loopback	Whether SONET/SDH loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Mode	Whether C-bit parity mode or M13 mode is enabled.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Output Field	Output Field Description
Keepalive statistics	Information about keepalive packets. ■ Input—Number of keepalive packets received by PPP. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*. ■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces.

Output Field	Output Field Description
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
DS3 alarms and defects	T3 media-specific defects 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. 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. ■ AIS—Alarm indicator signal. ■ EXZ—Excessive zeros. ■ FERF—Far-end failure. ■ IDLE—Idle alarm. ■ LCV—Line code violation. ■ LOF—Loss of frame. ■ LOS—Loss of signal. ■ PLL—Phase-locked loop out of lock. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.
SONET alarms and SONET defects	POS media-specific defects 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. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see Chapter 6 in GR-253, *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*. ■ AIS—Alarm indication signal ■ AIS-L—Alarm indication signal (line) ■ AIS-P—Alarm indication signal (path) ■ BERR—Bit error rate ■ BERR-SD—Bit error rate defect—signal degrade ■ BERR-SF—Bit error rate fault—signal fail ■ LOF—Loss of frame ■ LOL—Loss of light ■ LOP-P—Loss of pointer (path) ■ LOS—Loss of signal ■ PLL—Phase locked loop ■ PLM-P—Payload label mismatch ■ RDI—Remote defect indication ■ RDI-L—Remote defect indication (line) ■ RDI-P—Remote defect indication (path) ■ UNEQ-P—Unequipped (path)

Output Field	Output Field Description

SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Active alarms and defects, plus counts of specific SONET/SDH errors with detailed information.
	■ Seconds—Number of seconds the defect has been active.
	■ Count—Number of times that the defect has gone from inactive to active.
	■ State—State of the error. State other than OK indicates a problem.
	In addition to those listed under active alarms and defects, the SONET PHY, SONET section, SONET line, and SONET path fields report the following values:
	■ BIP—Bit Interleaved Parity
	  ■ BIP-B1—Bit Interleaved Parity for SONET/SDH section overhead
	  ■ BIP-B2—Bit Interleaved Parity for SONET/SDH line overhead
	  ■ BIP-B3—Bit Interleaved Parity for SONET/SDH path overhead
	■ ES—Errored Seconds
	  ■ ES-L—Errored Seconds (near-end line)
	  ■ ES-LFE—Errored Seconds (far-end line)
	  ■ ES-P—Errored Seconds (near-end STS path)
	  ■ ES-PFE—Errored Seconds (far-end STS path)
	  ■ ES-S—Errored Seconds (section)
	■ PHY Light—Loss of optical signal
	■ REI—Remote Error Indication
	  ■ REI-L—Remote Error Indication (line)
	  ■ REI-P—Remote Error Indication (path)
	■ SEF—Severely Errored Frame
	■ SEFS-S—Severely Errored Framing Seconds (section)
	■ SES—Severely Errored Seconds
	  ■ SES-L—Severely Errored Seconds (near-end line)
	  ■ SES-LFE—Severely Errored Seconds (far-end line)
	  ■ SES-P—Severely Errored Seconds (near-end STS path)
	  ■ SES-PFE—Severely Errored Seconds (far-end STS path)
	  ■ SES-S—Severely Errored Seconds (section)
	■ UAS—Unavailable Seconds
	  ■ UAS-L—Unavailable Seconds (near-end line)
	  ■ AS-LFE—Unavailable Seconds (far-end line)
	  ■ UAS-P—Unavailable Seconds (near-end STS path)
	  ■ UAS-PFE—Unavailable Seconds (far-end STS path)

DS3 media	(Extensive output only) Counts of T3 media-specific errors.
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Interface transmit queues	Name of the transmit queues and their associated statistics for each DS-3 channel on the Channelized OC-12 PIC. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.

Output Field	Output Field Description
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
DS3 BERT, DS-1 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see "test interface bert-start" on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT. It can be Pseudorandom, Repetitive, or All ones.
SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Counts of specific SONET/SDH errors.
Received SONET overhead, Transmitted SONET overhead	(Extensive output only) Values in the received and transmitted SONET/SDH overhead.
Received path trace, Transmitted path trace	(Extensive output only) Channelized OC-12 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. The received path trace value is the message received from the router at the other end of the fiber, and the transmitted path trace value is the message that this router transmits. This information is specific to each of the twelve Channelized OC-12 interfaces.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Output Field	Output Field Description
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 56 (00:00:02 ago), Output: 56 (00:00:06 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 16:59:03 PDT (18:23:36 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None

  Logical interface t3-0/3/0:0.0 (Index 11) (SNMP ifIndex 268)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 22.22.22.1, Local: 22.22.22.2
```

show interfaces brief (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0 brief
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 57 (00:00:04 ago), Output: 57 (00:00:07 ago)

  Logical interface t3-0/3/0:0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  22.22.22.2       --> 22.22.22.1
```

show interfaces detail (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0 detail
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21, Generation: 2719
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 57 (last seen 00:00:10 ago)
    Output: 58 (last sent 00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 16:59:03 PDT (18:23:53 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 1688                   88 bps
   Output bytes  :                 1702                   88 bps
   Input  packets:                  122                    0 pps
   Output packets:                  123                    0 pps
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None

  Logical interface t3-0/3/0:0.0 (Index 11) (SNMP ifIndex 268) (Generation 499)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 578, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 22.22.22.1, Local: 22.22.22.2, Broadcast: Unspecified,
        Generation: 981
```

show interfaces extensive (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0 extensive
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21, Generation: 2719
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 58 (last seen 00:00:04 ago)
    Output: 58 (last sent 00:00:07 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 16:59:03 PDT (18:23:58 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 1700                    0 bps
   Output bytes  :                 1714                    0 bps
   Input  packets:                  123                    0 pps
   Output packets:                  124                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 1100817, Bucket drops: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None
  DS3 media:            Seconds        Count  State
    AIS                       0            0  OK
    LOF                      18            1  OK
    LOS                       0            0  OK
    IDLE                      0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                      36       122399
    CCV                      72        91948
    LES                       0
    PES                      18
    PSES                     18
    CES                      18
    CSES                     18
    SEFS                     18
    UAS                       0
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
```

```
DS-3 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Algorithm: Unknown (0), Induced error rate: 10e-0
Interface transmit queues:
                 B/W  WRR       Packets        Bytes
  Queue0           0    0
    Transmitted:                0            0
    Drops:                0            0
    Errors:                0
    Queue1           0    0
    Transmitted:                0            0
    Drops:                0            0
    Errors:                0
  Queue2           0    0
    Transmitted:                0            0
    Drops:                0            0
    Errors:                0
  Queue3           0    0
    Transmitted:              124         1714
    Drops:                0            0
    Errors:                0
SONET PHY:            Seconds        Count  State
  PLL Lock                 0            0  OK
  PHY Light               20            1  OK
SONET section:
  BIP-B1                   0            0
  SEF                     20            1  OK
  LOS                     20            1  OK
  LOF                     20            1  OK
  ES-S                    20
  SES-S                   20
  SEFS-S                  20
SONET line:
  BIP-B2                   0            0
  REI-L                    0            0
  RDI-L                    0            0  OK
  AIS-L                    0            0  OK
  BERR-SF                 18            1  OK
  BERR-SD                  2            1  OK
  ES-L                    20
  SES-L                   20
  UAS-L                   10
  ES-LFE                   0
  SES-LFE                  0
  UAS-LFE                  0
SONET path:
  BIP-B3                   0            0
  REI-P                    0            0
  LOP-P                   20            1  OK
  AIS-P                    0            0  OK
  RDI-P                    0            0  OK
  UNEQ-P                   0            0  OK
  PLM-P                   20            1  OK
  ES-P                    20
  SES-P                   20
  UAS-P                   10
  ES-PFE                   0
  SES-PFE                  0
  UAS-PFE                  0
```

```
  Received SONET overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x04, C2(cmp) : 0x04, F2      : 0x00
    Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x00
    V5(cmp) : 0x00
  Transmitted SONET overhead:
    F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x04, F2      : 0x00, Z3      : 0x00
    Z4      : 0x00, V5      : 0x00
  Received path trace: t3-0/3/0:0
    74 33 2d 30 2f 33 2f 30 3a 30 00 00 00 00 0d 0a   t3-0/3/0:0......
  Transmitted path trace: t3-0/3/0:0
    74 33 2d 30 2f 33 2f 30 3a 30 00 00 00 00 00 00   t3-0/3/0:0......
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                             %          bps   %        bytes
    0 best-effort            95     42499200  95            0      low    none
    3 network-control         5      2236800   5            0      low    none

  Logical interface t3-0/3/0:0.0 (Index 11) (SNMP ifIndex 268) (Generation 499)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 578, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 22.22.22.1, Local: 22.22.22.2, Broadcast: Unspecified,
        Generation: 981
```

show interfaces media (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0 media
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 58 (00:00:10 ago), Output: 59 (00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 16:59:03 PDT (18:24:04 ago)
  Input rate     : 88 bps (0 pps)
  Output rate    : 88 bps (0 pps)
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None
  DS3 errors:
    BPV: 0, EXZ: 0, LCV: 0, PCV: 122399, CCV: 91948
```

```
  Interface transmit queues:
                  B/W  WRR       Packets        Bytes
    Queue0          0    0
      Transmitted:                0            0
      Drops:                0            0
      Errors:                0
    Queue1          0    0
      Transmitted:                0            0
      Drops:                0            0
      Errors:                0
    Queue2          0    0
      Transmitted:                0            0
      Drops:                0            0
      Errors:                0
    Queue3          0    0
      Transmitted:              125         1726
      Drops:                0            0
      Errors:                0
  SONET errors:
    BIP-B1: 0, BIP-B2: 0, REI-L: 0, BIP-B3: 0, REI-P: 0, BIP-BIP2: 0
  Received path trace: t3-0/3/0:0
  Transmitted path trace: t3-0/3/0:0
```

show interfaces statistics (for Channelized OC-12 Interfaces)

```
user@host> show interfaces t3-0/3/0:0 statistics
Physical interface: t3-0/3/0:0, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 21
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: T3,
  Loopback: None, SONET Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 59 (00:00:06 ago), Output: 59 (00:00:07 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 16:59:03 PDT (18:24:11 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 1100817, Output errors: 0
  DS3   alarms   : None
  SONET alarms   : None
  DS3   defects  : None
  SONET defects  : None

  Logical interface t3-0/3/0:0.0 (Index 11) (SNMP ifIndex 268)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 22.22.22.1, Local: 22.22.22.2
```

Chapter 14

Channelized STM-1 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Channelized STM-1 interfaces.

show interfaces (for Channelized STM-1 Interfaces)

Syntax

show interfaces e1-*fpc/pic/port:channel* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces e1-*fpc/pic/port:channel* <brief | detail | extensive> <media>

show interfaces e1-*fpc/pic/port:channel* <brief | detail | extensive> <statistics>

Description

Display status information about Channelized STM-1 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

e1-*fpc/pic/port:channel*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Channelized STM-1 Interfaces) on page 292
show interfaces brief (for Channelized STM-1 Interfaces) on page 293
show interfaces detail (for Channelized STM-1 Interfaces) on page 293
show interfaces extensive (for Channelized STM-1 Interfaces) on page 294
show interfaces media (for Channelized STM-1 Interfaces) on page 296
show interfaces statistics (for Channelized STM-1 Interfaces) on page 297

Output Fields at a Glance

Table 27 summarizes which information is included in the output fields of each show interfaces command option for STM-1 interfaces. In this table, output fields are listed in alphabetical order. In Table 28, the output fields are listed in more detail in the order in which they are displayed.

Table 27: Channelized STM-1 Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—E1 media-specific defects that can render the interface unable to pass packets.
Detail Extensive	BERT time period—Configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16 bit.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Device Flags Field—Information about the physical device.
Extensive	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.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—BERT (Bit Error Rate Testing) checks the quality of the line.
Extensive	E1 media—Counts of E1 media-specific errors.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	Framing—Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—Instance number of the interface.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.

Options	Field Description
Extensive	Interface transmit queues—Transmit queue statistics and configuration for each E1 channel on the Channelized STM-1 to E1 PIC.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface. The default MTU size is 1504.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
All	SONET alarms and SONET defects—POS media-specific defects 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.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 28: Channelized STM-1 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Output Field	Output Field Description
Keepalive statistics	Information about keepalive packets. ■ Input—Number of keepalive packets received by PPP. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*. ■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects	E1 media-specific defects 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. 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. ■ LOS—Loss of signal. ■ LOF—Loss of frame. ■ AIS—Alarm indication signal. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.

Output Field	Output Field Description
E1 media	(Extensive output only) Active alarms and defects, plus counts of specific E1 errors with detailed information. ■ Seconds—Number of seconds the defect has been active. ■ Count—Number of times that the defect has gone from inactive to active. ■ State—State of the error. State other than OK indicates a problem. The E1 media specific error types can be: ■ AIS—Alarm indication signal ■ BEE—Bit error ■ BES—Bit error seconds ■ BPV—Bipolar violation ■ CS—Carrier state ■ ES—Errored Seconds ■ EXZ—Excessive zeros ■ FEBE—Far-end block error ■ LCV—Line code violation ■ LES—Line error seconds ■ LOF—Loss of frame ■ LOS—Loss of signal ■ PCV—Pulse code violation ■ SEF—Severely Errored Frame ■ SEFS-S—Severely Errored Framing Seconds (section) ■ SES—Severely Errored Seconds ■ UAS—Unavailable Seconds ■ YELLOW—Errors at the remote site receiver
SONET alarms and SONET defects	POS media-specific defects 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. 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see Chapter 6 in GR-253, *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*. ■ AIS—Alarm indication signal ■ AIS-L—Alarm indication signal (line) ■ AIS-P—Alarm indication signal (path) ■ BERR—Bit error rate ■ BERR-SD—Bit error rate defect–signal degrade ■ BERR-SF—Bit error rate fault–signal fail ■ LOF—Loss of frame ■ LOL—Loss of light ■ LOP-P—Loss of pointer (path) ■ LOS—Loss of signal ■ PLL—Phase locked loop ■ PLM-P—Payload label mismatch ■ RDI—Remote defect indication ■ RDI-L—Remote defect indication (line) ■ RDI-P—Remote defect indication (path) ■ UNEQ-P—Unequipped (path)

Output Field	Output Field Description

SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Active alarms and defects, plus counts of specific SONET/SDH errors with detailed information.
	■ Seconds—Number of seconds the defect has been active.
	■ Count—Number of times that the defect has gone from inactive to active.
	■ State—Indicates state of the error. State other than OK indicates a problem.
	In addition to those listed under Active alarms and Active defects, the SONET PHY, SONET section, SONET line, and SONET path fields report the following values:
	■ BIP—Bit Interleaved Parity
	  ■ BIP-B1—Bit Interleaved Parity for SONET/SDH section overhead
	  ■ BIP-B2—Bit Interleaved Parity for SONET/SDH line overhead
	  ■ BIP-B3—Bit Interleaved Parity for SONET/SDH path overhead
	■ ES—Errored Seconds
	  ■ ES-L—Errored Seconds (near-end line)
	  ■ ES-LFE—Errored Seconds (far-end line)
	  ■ ES-P—Errored Seconds (near-end STS path)
	  ■ ES-PFE—Errored Seconds (far-end STS path)
	  ■ ES-S—Errored Seconds (section)
	■ PHY Light—Loss of optical signal
	■ REI—Remote Error Indication
	  ■ REI-L—Remote Error Indication (line)
	  ■ REI-P—Remote Error Indication (path)
	■ SEF—Severely Errored Frame
	■ SEFS-S—Severely Errored Framing Seconds (section)
	■ SES—Severely Errored Seconds
	  ■ SES-L—Severely Errored Seconds (near-end line)
	  ■ SES-LFE—Severely Errored Seconds (far-end line)
	  ■ SES-P—Severely Errored Seconds (near-end STS path)
	  ■ SES-PFE—Severely Errored Seconds (far-end STS path)
	  ■ SES-S—Severely Errored Seconds (section)
	■ UAS—Unavailable Seconds
	  ■ UAS-L—Unavailable Seconds (near-end line)
	  ■ AS-LFE—Unavailable Seconds (far-end line)
	  ■ UAS-P—Unavailable Seconds (near-end STS path)
	  ■ UAS-PFE— Unavailable Seconds (far-end STS path)

HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Interface transmit queues	Names of the transmit queues and their associated statistics for each E1 channel on the Channelized STM-1 to E1 PIC. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.

Output Field	Output Field Description
Buildout	Buildout setting.
DS3 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see "test interface bert-start" on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT.
SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Counts of specific SONET/SDH errors.
Received SONET overhead, Transmitted SONET overhead	(Extensive output only) Values in the received and transmitted SONET/SDH overhead.
Received path trace, Transmitted path trace	(Extensive output only) Channelized STM-1 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. The received path trace value is the message received from the router at the other end of the fiber, and the transmitted path trace value is the message that this router transmits. This information is specific to each of the twelve Channelized STM-1 interfaces.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Output Field	Output Field Description
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in “Family Flags Field” on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in “Address Flags Field” on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 18 (00:00:05 ago), Output: 18 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:02:59 PDT (17:23:22 ago)
  Input rate     : 48 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  DS1   alarms   : None
  SDH   alarms   : None
  DS1   defects  : None
  SDH   defects  : None

  Logical interface e1-1/0/0:1.0 (Index 10) (SNMP ifIndex 369)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 19.19.19.3, Local: 19.19.19.4
```

show interfaces brief (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1 brief
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 18 (00:00:11 ago), Output: 18 (00:00:07 ago)

  Logical interface e1-1/0/0:1.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  19.19.19.4       --> 19.19.19.3
```

show interfaces detail (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1 detail
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285, Generation: 2915
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 20 (last seen 00:00:02 ago)
    Output: 19 (last sent 00:00:10 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:02:59 PDT (17:23:40 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  579                   48 bps
   Output bytes  :                  630                   48 bps
   Input  packets:                   45                    0 pps
   Output packets:                   45                    0 pps
  DS1   alarms   : None
  SDH   alarms   : None
  DS1   defects  : None
  SDH   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)

  Logical interface e1-1/0/0:1.0 (Index 10) (SNMP ifIndex 369) (Generation 496)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 575, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 19.19.19.3, Local: 19.19.19.4, Broadcast: Unspecified,
        Generation: 975
```

show interfaces extensive (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1 extensive
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285, Generation: 2915
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 20 (last seen 00:00:07 ago)
    Output: 20 (last sent 00:00:03 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:02:59 PDT (17:23:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  592                   48 bps
   Output bytes  :                  644                   48 bps
   Input  packets:                   46                    0 pps
   Output packets:                   46                    0 pps
  Input errors:
    Errors: 0, Drops: 9, Framing errors: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 11, L2 mismatch timeouts: 0,
    HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  SDH   alarms   : None
  DS1   defects  : None
  SDH   defects  : None
  E1  media:            Seconds        Count  State
    SEF                       0            0  OK
    BEE                       0            0  OK
    AIS                     124            1  OK
    LOF                     124            1  OK
    LOS                       0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CS                        0            0
    FEBE                      0            0
    LES                     124
    ES                      125
    SES                     124
    SEFS                    124
    BES                       0
    UAS                      37
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5           529         6348            0            0
```

```
HDLC configuration:
  Giant threshold: 0, Runt threshold: 0
  Timeslots      : All active
  Line encoding: HDB3
DS1 BERT configuration:
  BERT time period: 10 seconds, Elapsed: 0 seconds
  Induced Error rate: 10e-0, Algorithm: 2^15 - 1, O.151, Pseudorandom (9)
SDH PHY:              Seconds        Count  State
  PLL Lock                  0            0  OK
  PHY Light                 0            0  OK
SDH regenerator section:
  RS-BIP8                   0            0
  OOF                     125            1  OK
  LOS                     125            1  OK
  LOF                     125            1  OK
  RS-ES                   125
  RS-SES                  125
  RS-SEFS                 125
SDH multiplex section:
  MS-BIP24                  0            0
  MS-FEBE                   0            0
  MS-FERF                   0            0  OK
  MS-AIS                  125            1  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  MS-ES                   125
  MS-SES                  125
  MS-UAS                  115
  MS-ES-FE                  0
  MS-SES-FE                 0
  MS-UAS-FE                 0
SDH path:
  HP-BIP8                   0            0
  HP-FEBE                   0            0
  HP-LOP                    0            0  OK
  HP-AIS                  125            1  OK
  HP-FERF                   0            0  OK
  HP-UNEQ                   0            0  OK
  HP-PLM                  125            1  OK
  HP-ES                   125
  HP-SES                  125
  HP-UAS                  115
  HP-ES-FE                  0
  HP-SES-FE                 0
  HP-UAS-FE                 0
SDH tu:
  TU-BIP2                   0            0
  TU-FEBE                 124            1
  TU-LOP                    0            0  OK
  TU-AIS                  124            1  OK
  TU-FERF                 124            1  OK
  TU-UNEQ                   0            0  OK
  TU-PLM                  124            1  OK
  TU-ES                   125
  TU-SES                  125
  TU-UAS                  115
  TU-ES-FE                  0
  TU-SES-FE                 0
  TU-UAS-FE                 0
```

```
  Received SDH overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x02, C2(cmp) : 0x02, F2      : 0x00
    Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x02
    V5(cmp) : 0x02
  Transmitted SDH overhead:
    F1      : 0x00, J0      : 0x00, K1      : 0x00, K2      : 0x00
    S1      : 0x00, C2      : 0x02, F2      : 0x00, Z3      : 0x00
    Z4      : 0x00, V5      : 0x02
  Received path trace:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  Transmitted path trace:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  Packet Forwarding Engine configuration:
    Destination slot: 1, PLP byte: 2 (0x07)
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                              %          bps   %        bytes
    0 best-effort            95      1945600  95            0      low    none
    3 network-control         5       102400   5            0      low    none

  Logical interface e1-1/0/0:1.0 (Index 10) (SNMP ifIndex 369) (Generation 496)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500, Generation: 575, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 19.19.19.3, Local: 19.19.19.4, Broadcast: Unspecified,
        Generation: 975
```

show interfaces media (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1 media
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 21 (00:00:07 ago), Output: 21 (00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:02:59 PDT (17:23:53 ago)
  Input rate     : 48 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  DS1   alarms   : None
  SDH   alarms   : None
  DS1   defects  : None
  SDH   defects  : None
  E1 errors:
    BPV: 0, EXZ: 0, LCV: 0, PCV: 0, CS: 0, FEBE: 0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95             0            0            0            0
    Queue1      5    5           576         6912            0            0
  SDH errors:
    RS-BIP8: 0, MS-BIP24: 0, MS-FEBE: 0, HP-BIP8: 0, HP-FERF: 0, TU-BIP2: 0
  Received path trace:
  Transmitted path trace:
```

show interfaces statistics (for Channelized STM-1 Interfaces)

```
user@host> show interfaces e1-1/0/0:1 statistics
Physical interface: e1-1/0/0:1, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 285
  Link-level type: PPP, MTU: 1504, SDH mode, Speed: E1, Loopback: None, FCS: 16,
  Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 22 (00:00:03 ago), Output: 22 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-23 17:02:59 PDT (17:24:00 ago)
  Statistics last cleared: Never
  Input rate     : 48 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  Input errors: 20, Output errors: 0
  DS1   alarms   : None
  SDH   alarms   : None
  DS1   defects  : None
  SDH   defects  : None

  Logical interface e1-1/0/0:1.0 (Index 10) (SNMP ifIndex 369)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 19.19.19.3, Local: 19.19.19.4
```

Chapter 15

E-1 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot E1 interfaces.

show interfaces (for E1 Interfaces)

Syntax

show interfaces e1-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces e1-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces e1-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about E1 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

e1-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for E1 interfaces) on page 307
show interfaces brief (for E1 interfaces) on page 308
show interfaces detail (for E1 interfaces) on page 308
show interfaces extensive (for E1 interfaces) on page 309
show interfaces media (for E1 interfaces) on page 311
show interfaces statistics (for E1 interfaces) on page 312

Output Fields at a Glance

Table 29 summarizes which information is included in the output fields of each show interfaces command option for E1 interfaces. In this table, output fields are listed in alphabetical order. In Table 30, the output fields are listed in more detail in the order in which they are displayed.

Table 29: E1 Show Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—E1 media-specific defects that can render the interface unable to pass packets.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
All	Buildout—Buildout setting.
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Extensive	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.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Test.
Detail Extensive	DS3 BERT configuration—The BERT (Bit Error Rate Test) checks the quality of the line.
All	Device Flags Field—Information about the physical device.
Extensive	E1 media—Counts of E1 media-specific errors.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled—State of the interface.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
All	Framing—Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Giants—Frames received that are larger than the giant threshold.
	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.

Options	Field Description
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	Interface transmit queues—Name of the transmit queues and their associated statistics.
All	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link-level type—Encapsulation being used on the physical interface.
All	Link Flags Field—Information about the link.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface. The default MTU size is 1504.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Detail Extensive	Protocol running on the logical interface, such as iso, inet6, mpls.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 30: E1 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Output Field	Output Field Description
Keepalive statistics	Information about keepalive packets. ■ Input—Number of keepalive packets received by PPP. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*. ■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects	E1 media-specific defects 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. 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. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see *Bellcore Telcordia GR-499-CORE*. ■ LOS—Loss of signal. ■ LOF—Loss of frame. ■ AIS—Alarm indication signal. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.

Output Field	Output Field Description
E1 media	(Extensive output only) Active alarms and defects, plus counts of specific E1 errors with detailed information. ■ Seconds—Number of seconds the defect has been active. ■ Count—Number of times that the defect has gone from inactive to active. ■ State—State of the error. State other than OK indicates a problem. The E1 media specific error types can be: ■ AIS—Alarm indication signal ■ BEE—Bit error ■ BES—Bit error seconds ■ BPV—Bipolar violation ■ CS—Carrier state ■ ES—Errored Seconds ■ EXZ—Excessive zeros ■ FEBE—Far-end block error ■ LCV—Line code violation ■ LES—Line error seconds ■ LOF—Loss of frame ■ LOS—Loss of signal ■ PCV—Pulse code violation ■ SEF—Severely Errored Frame ■ SEFS-S—Severely Errored Framing Seconds (section) ■ SES—Severely Errored Seconds ■ UAS—Unavailable Seconds ■ YELLOW—Errors at the remote site receiver
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Interface transmit queues	Names of the transmit queues and their associated statistics. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
Buildout	Buildout setting.

Output Field	Output Field Description
DS1 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see "test interface bert-start" on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Tail queue drops	Number of packets dropped due to bandwidth constraints. Indicates that packets are queued to send out at a rate faster than allowed.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name,* Output: *output-filter-name.*

Output Field	Output Field Description
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in “Address Flags Field” on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for E1 interfaces)

```
user@host> show interfaces  e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Interface index: 100, SNMP ifIndex: 243
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 99 (00:00:05 ago), Output: 100 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 10:52:40 PDT (04:33:20 ago)
  Input rate     : 404792 bps (48 pps)
  Output rate    : 844632 bps (102 pps)
  DS1   alarms   : None
  DS1   defects  : None

  Logical interface e1-4/3/0.0 (Index 200) (SNMP ifIndex 253)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc
    Protocol mpls, MTU: 1488
      Flags: None
```

show interfaces brief (for E1 interfaces)

```
user@host> show interfaces brief e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 99 (00:00:05 ago), Output: 100 (00:00:01 ago)

  Logical interface e1-4/3/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  10.41.251.1/24
    iso
    inet6 2060::251:1/126
          fe80::2a0:a5ff:fe12:7dc/64
    mpls
```

show interfaces detail (for E1 interfaces)

```
user@host> show interfaces detail e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Interface index: 100, SNMP ifIndex: 243, Generation: 1143
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 99 (last seen 00:00:05 ago)
    Output: 100 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 10:52:40 PDT (04:33:20 ago)
  Statistics last cleared: 2002-04-26 15:09:06 PDT (00:16:54 ago)
  Traffic statistics:
   Input  bytes  :              8497524               404792 bps
   Output bytes  :             20954183               844632 bps
   Input  packets:                 8949                   48 pps
   Output packets:                20734                  102 pps
  DS1   alarms   : None
  DS1   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^11 - 1, O.152 and O.153 (2047 type),
Pseudorandom (8)
```

```
  Logical interface e1-4/3/0.0 (Index 200) (SNMP ifIndex 253) (Generation 1236)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
    Generation: 3697 Route table: 0
      Filters: Input: sample-all
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1, Broadcast: Unspecified,
        Generation: 6156
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500Generation: 3703 Route table: 0
      Flags: None
    Generation: 3704 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1,
        Broadcast: Unspecified, Generation: 6163
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc,
        Broadcast: Unspecified, Generation: 6164
    Protocol mpls, MTU: 1488
      Flags: None
    Generation: 3705 Route table: 0
```

show interfaces extensive (for E1 interfaces)

```
user@host> show interfaces extensive e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Interface index: 100, SNMP ifIndex: 243, Generation: 1143
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 99 (last seen 00:00:05 ago)
    Output: 100 (last sent 00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 10:52:40 PDT (04:33:20 ago)
  Statistics last cleared: 2002-04-26 15:09:06 PDT (00:16:54 ago)
  Traffic statistics:
   Input  bytes  :              8497524                404792 bps
   Output bytes  :             20954183                844632 bps
   Input  packets:                 8949                    48 pps
   Output packets:                20734                   102 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  DS1   defects  : None
```

```
  E1  media:            Seconds        Count  State
    SEF                       0            0  OK
    BEE                       0            0  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CS                        0            0
    FEBE                      0            0
    LES                       0
    ES                        0
    SES                       0
    SEFS                      0
    BES                       0
    UAS                       0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95          3627      3743064            0            0
    Queue1      5    5         16959     17039685            0            0
  HDLC configuration:
    Giant threshold: 1514, Runt threshold: 0
    Timeslots      : All active
    Line encoding: HDB3
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: 2^11 - 1, O.152 and O.153 (2047 type),
Pseudorandom (8)
  Packet Forwarding Engine configuration:
    Destination slot: 4, PLP byte: 2 (0x00)
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                              %          bps   %        bytes
    0 best-effort            95      1945600  95            0      low    none
    3 network-control         5       102400   5            0      low    none

  Logical interface e1-4/3/0.0 (Index 200) (SNMP ifIndex 253) (Generation 1236)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
    Generation: 3697 Route table: 0
      Filters: Input: sample-all
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1, Broadcast: Unspecified,
        Generation: 6156
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500Generation: 3703 Route table: 0
      Flags: None
    Generation: 3704 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1,
        Broadcast: Unspecified, Generation: 6163
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc,
        Broadcast: Unspecified, Generation: 6164
    Protocol mpls, MTU: 1488
      Flags: None
    Generation: 3705 Route table: 0
```

show interfaces media (for E1 interfaces)

```
user@host> show interfaces media e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Interface index: 100, SNMP ifIndex: 243
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 99 (00:00:05 ago), Output: 100 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 10:52:40 PDT (04:33:20 ago)
  Input rate     : 404792 bps (48 pps)
  Output rate    : 844632 bps (102 pps)
  DS1   alarms   : None
  DS1   defects  : None
  E1 errors:
    BPV: 0, EXZ: 0, LCV: 0, PCV: 0, CS: 0, FEBE: 0
  Interface transmit queues:
              B/W  WRR       Packets        Bytes        Drops       Errors
    Queue0     95   95          3627      3743064            0            0
    Queue1      5    5         16959     17039685            0            0
```

show interfaces statistics (for E1 interfaces)

```
user@host> show interfaces statistics e1-4/3/0
Physical interface: e1-4/3/0, Enabled, Physical link is Up
  Interface index: 100, SNMP ifIndex: 243
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: E1,
  Loopback: None, FCS: 16, Framing: G704
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 99 (00:00:05 ago), Output: 100 (00:00:01 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Opened, mpls: Opened
  CHAP state: Not-configured
  Last flapped   : 2002-04-26 10:52:40 PDT (04:33:20 ago)
  Statistics last cleared: 2002-04-26 15:09:06 PDT (00:16:54 ago)
  Input rate     : 404792 bps (48 pps)
  Output rate    : 844632 bps (102 pps)
  Input errors: 0, Output errors: 0
  DS1   alarms   : None
  DS1   defects  : None

  Logical interface e1-4/3/0.0 (Index 200) (SNMP ifIndex 253)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Bandwidth: 0
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.41.251/24, Local: 10.41.251.1
    Protocol iso, MTU: 1500
      Flags: None
    Protocol inet6, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 2060::251:0/126, Local: 2060::251:1
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe12:7dc
    Protocol mpls, MTU: 1488
      Flags: None
```

Chapter 16
E-3 Interfaces Monitoring and Troubleshooting

This chapter displays the command-line interface (CLI) commmands you can use to monitor and troubleshoot E3 interfaces.

show interfaces (for E3 Interfaces)

Syntax

show interfaces e3-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces e3-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces e3-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about E3 router interfaces.

Options

none—Display information about all interfaces.

brief—Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

e3-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for E3 interfaces) on page 320
show interfaces brief (for E3 interfaces) on page 320
show interfaces detail (for E3 interfaces) on page 320
show interfaces extensive (for E3 interfaces) on page 321
show interfaces media (for E3 interfaces) on page 322
show interfaces statistics (for E3 interfaces) on page 323

Output Fields at a Glance

Table 31 summarizes which information is included in the output fields of each show interfaces command option for E3 interfaces. In this table, output fields are listed in alphabetical order. In Table 32, the output fields are listed in more detail in the order in which they are displayed.

Table 31: E3 Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—E3 media-specific defects that can render the interface unable to pass packets.
All	Clocking—Reference clock source. It can be Internal or External.
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
All	Buildout—Buildout setting.
All	FCS—Frame check sequence on the interface (either 16 or 32).
All	Device Flags Field—Information about the physical device.
Extensive	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.
Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—The BERT (Bit Error Rate Test) checks the quality of the line.
Detail Extensive	DS3 media—Counts of E3 media-specific errors.
Extensive	DS3 media—Counts of E3 media-specific errors.
Extensive	DSU configuration—Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	Framing—Physical layer framing format used on the link.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
All	Keepalive settings—Configured settings for keepalives.
All	Link Flags Field—Information about the link.

Options	Field Description
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
Logical Interface	
All	Addresses flags—Information about the address.
All	Addresses—Addresses associated with the logical interface.
Detail Extensive	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
All	Encapsulation—Encapsulation on the logical interface.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
Extensive	ITU LMI settings—Settings for Link Management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Extensive	LMI Statistics—Statistics about Link Management, including a count of packets sent and received, and the time of the last activity.
All	Local—IP address of the logical interface.
All	Logical interface flags—Information about the logical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number, and the logical interface's SNMP interface index number.
All	MTU—MTU size on the logical interface. The default MTU size is 4474.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail Extensive	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Table 32: E3 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.

Output Field	Output Field Description
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in “Device Flags Field” on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in “Link Flags Field” on page 191.
LMI settings	(Extensive output only) Settings for Link Management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the Link Management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link’s status from down to up. The range is 1 through 255, with a default of 1.

Keepalive statistics	Information about keepalive packets.
	■ Input—Number of keepalive packets received by PPP.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*.
	■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.

LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.

Output Field	Output Field Description
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with Radius authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.

Output Field	Output Field Description
Active alarms and Active defects	Active alarms and Active defects—E3 media-specific defects 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. 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. AIS—Alarm indication signal EXZ—Excessive zeros FERF—For end receive failures IDLE— Idle code detected LCD—Loss of cell delineation LCV—Line code violation LOF—Loss of frame LOS—Loss of signal PLL—Phase locked loop YLW—Remote defect indication
E3 media	(Extensive output only) Counts of E3 media-specific errors.
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Line encoding	Line encoding used. It is always HDB3.
Interface transmit queues	Name of the transmit queues and their associated statistics. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. Scrambling—Payload scrambling. It can be Enabled or Disabled. Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
Buildout	Buildout setting.
DS3 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see “test interface bert-start” on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT. It can be Pseudorandom, Repetitive, or All ones.

Output Field	Output Field Description

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only non-zero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for E3 interfaces)

```
user@host> show interfaces  e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Interface index: 77, SNMP ifIndex: 840
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 49 (00:00:04 ago), Output: 47 (00:00:10 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:29 PDT (00:09:36 ago)
  Input rate     : 264 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface e3-0/3/0.0 (Index 8) (SNMP ifIndex 856)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.20.115.0/30, Local: 10.20.115.2
```

show interfaces brief (for E3 interfaces)

```
user@host> show interfaces brief e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 49 (00:00:04 ago), Output: 47 (00:00:10 ago)

  Logical interface e3-0/3/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  10.20.115.2/30
```

show interfaces detail (for E3 interfaces)

```
user@host> show interfaces detail e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Interface index: 77, SNMP ifIndex: 840, Generation: 76
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 49 (last seen 00:00:04 ago)
    Output: 47 (last sent 00:00:10 ago)
```

```
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:29 PDT (00:09:36 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:08:08 ago)
  Traffic statistics:
   Input  bytes  :                 8800                  264 bps
   Output bytes  :                 9568                    0 bps
   Input  packets:                  192                    0 pps
   Output packets:                  192                    0 pps
  Active alarms  : None
  Active defects : None

  Logical interface e3-0/3/0.0 (Index 8) (SNMP ifIndex 856) (Generation 8)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 16 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.20.115.0/30, Local: 10.20.115.2, Broadcast: Unspecified,
        Generation: 18
```

show interfaces extensive (for E3 interfaces)

```
user@host> show interfaces extensive e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Interface index: 77, SNMP ifIndex: 840, Generation: 76
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 49 (last seen 00:00:04 ago)
    Output: 47 (last sent 00:00:10 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:29 PDT (00:09:36 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:08:08 ago)
  Traffic statistics:
   Input  bytes  :                 8800                  264 bps
   Output bytes  :                 9568                    0 bps
   Input  packets:                  192                    0 pps
   Output packets:                  192                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 1, Bucket drops: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0
  Active alarms  : None
  Active defects : None
  E3  media:            Seconds        Count  State
    PLL Lock                  0            0  OK
    Reframing                 0            0  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
```

```
    IDLE                      0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    LES                       0
    SEFS                      0
    UAS                       0
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  DSU configuration:
    Compatibility mode: Kentrox, Scrambling: Disabled
  E3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: 2^3 - 1, Pseudorandom (1), Induced Error rate: 10e-0
  Packet Forwarding Engine configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %          bps   %        bytes
    0 best-effort             0            0   0            0      low    none
    1 expedited-forwarding    0            0   0            0      low    none
    2 assured-forwarding      0            0   0            0      low    none
    3 network-control         0            0   0            0      low    none

  Logical interface e3-0/3/0.0 (Index 8) (SNMP ifIndex 856) (Generation 8)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 16 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.20.115.0/30, Local: 10.20.115.2, Broadcast: Unspecified,
        Generation: 18
```

show interfaces media (for E3 interfaces)

```
user@host> show interfaces media e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Interface index: 77, SNMP ifIndex: 840
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 49 (00:00:05 ago), Output: 47 (00:00:11 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:29 PDT (00:09:37 ago)
  Input rate     : 264 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  E3 errors:
    BPV: 0, EXZ: 0, LCV: 0
    PCV: 0
```

show interfaces statistics (for E3 interfaces)

```
user@host> show interfaces statistics e3-0/3/0
Physical interface: e3-0/3/0, Enabled, Physical link is Up
  Interface index: 77, SNMP ifIndex: 840
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: E3,
  Loopback: None, FCS: 16, Framing: G751
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 49 (00:00:05 ago), Output: 47 (00:00:11 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:29 PDT (00:09:37 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:08:09 ago)
  Input rate     : 264 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 1, Output errors: 0
  Active alarms  : None
  Active defects : None

  Logical interface e3-0/3/0.0 (Index 8) (SNMP ifIndex 856)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.20.115.0/30, Local: 10.20.115.2
```

Chapter 17
ES Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot ES interfaces.

show interfaces (for ES Interfaces)

Syntax

show interfaces es-*fpc/pic/port:es* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces es-*fpc/pic/port:es* <brief | detail | extensive> <media>

show interfaces es-*fpc/pic/port:es* <brief | detail | extensive> <statistics>

Description

Display status information about ES router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

es-*fpc/pic/port:es channel*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for ES interfaces) on page 330
show interfaces brief (for ES interfaces) on page 330
show interfaces detail (for ES interfaces) on page 331
show interfaces extensive (for ES interfaces) on page 332
show interfaces media (for ES interfaces) on page 332
show interfaces statistics (for ES interfaces) on page 333

Output Fields at a Glance

Table 33 summarizes which information is included in the output fields of each show interfaces command option for ES interfaces. In this table, output fields are listed in alphabetical order. In Table 34, the output fields are listed in more detail in the order in which they are displayed.

Table 33: ES Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
Detail	Anti-replay failures—Total number of anti-replay failures seen on all tunnels configured on the ES PIC.
Detail	Authentication—Total number of authentication failures seen on all tunnels configured on the ES PIC.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Enabled—State of the interface.
All	Flags—Information about the physical device and interface.
Detail	Generation—A unique number for use by Juniper Networks Customer Support only.
Detail	Hold-times—Hold time configured with the hold-time statement. The default is 90 seconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Detail	IPSec security association—Name of the security association.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link-level type—Describes the link layer type.
All	MTU—MTU size on the physical interface.
All	Physical interface—Name of the physical interface.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Standard Detail	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail	Total authentication failures—Total number of authentication failures seen on all logical tunnels configured on the ES PIC.
Detail	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
All	Type—Encapsulation being used on the interface.

Options	Field Description
Logical Interface	
Detail	Authentication failures—Number of authentication failures seen on a tunnel configured on the ES PIC.
Detail	Broadcast—Broadcast address on the logical interface.
Standard Detail	Destination—For a point-to-point link, the address of the remote side of the link. For multicast links, the network address.
All	Encapsulation—Encapsulation on the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Flags—Information about the logical interface.
Detail	IPSec Security association—Name of the security association.
Detail	Local—IP address of the logical interface.
Detail	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
Standard Detail	Protocol—Protocol running on the logical interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Detail	Receive sequence number—Receive sequence number of the replay window. This is applicable only if authentication is configured on the tunnel.
Detail	Route table—Name of route table.
Detail	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail	Transmit sequence number—Transmit sequence number of the replay window. This is applicable only if authentication is configured on the tunnel.

Table 34: ES Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Type	Encapsulation being used on the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).

Output Field	Output Field Description
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Anti-replay failures	Total number of anti-replay failures seen on all tunnels configured on the ES PIC.
Total Authentication failures	(Physical interface) Total number of authentication failures seen on all tunnels configured on the ES PIC.
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Generation	A unique number for use by Juniper Networks Customer Support only.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.
IPSec Security association	Name of the security association.
Authentication failures	(Logical interface) Number of authentication failures seen on a tunnel configured on the ES PIC.

Output Field	Output Field Description
Receive sequence number	Receive sequence number of the replay window. This is applicable only if authentication is configured on the tunnel.
Transmit sequence number	Transmit sequence number of the replay window. This is applicable only if authentication is configured on the tunnel.

show interfaces (standard) (for ES interfaces)

```
user@host> show interfaces es-0/2/0
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Interface index: 18, SNMP ifIndex: 22
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface es-0/2/0.0 (Index 6) (SNMP ifIndex 39)
    Flags: Point-To-Point SNMP-Traps Encapsulation: IPSec
  Input packets : 0
  Output packets: 0
    Protocol inet, IPSec Security association: sa2, MTU: 1400, Flags:
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 4.5.6.7, Local: 1.2.3.4
```

show interfaces brief (for ES interfaces)

```
user@host> show interfaces es-0/2/0 brief
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface es-0/2/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: IPSec
    inet  1.2.3.4          --> 4.5.6.7
```

show interfaces detail (for ES interfaces)

```
user@host> show interfaces es-0/2/0 detail
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Interface index: 15, SNMP ifIndex: 19, Generation: 78
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   Anti-replay failures     : 0
   Authenticaion failures   : 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

  Logical interface es-0/2/0.0 (Index 6) (SNMP ifIndex 21) (Generation 29)
    Flags: Point-To-Point SNMP-Traps Encapsulation: IPSec
    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
```

show interfaces extensive (for ES interfaces)

```
user@host> show interfaces es-0/2/0 extensive
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Interface index: 18, SNMP ifIndex: 22, Generation: 17
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
   Anti-replay failures       : 0
   Authentication failures    : 0

  Logical interface es-0/2/0.0 (Index 6) (SNMP ifIndex 39) (Generation 5)
    Flags: Point-To-Point SNMP-Traps Encapsulation: IPSec
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   0                    0 pps
     Output packets:                   0                    0 pps
    Protocol inet, IPSec Security association: sa2, MTU: 1400, Flags:
    Generation: 7 Route table: 0 Authentication failures: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 4.5.6.7, Local: 1.2.3.4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces media (for ES interfaces)

```
user@host> show interfaces es-0/2/0 media
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Interface index: 15, SNMP ifIndex: 19
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces statistics (for ES interfaces)

```
user@host> show interfaces es-0/2/0 statistics
Physical interface: es-0/2/0, Enabled, Physical link is Up
  Interface index: 18, SNMP ifIndex: 22
  Type: IPSec, Link-level type: IPSec-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Anti-replay failures       : 0
  Authentication failures    : 0

  Logical interface es-0/2/0.0 (Index 6) (SNMP ifIndex 39)
    Flags: Point-To-Point SNMP-Traps Encapsulation: IPSec
  Input packets : 0
  Output packets: 0
    Protocol inet, IPSec Security Association: sa2, MTU: 1400, Flags:
    Authentication failures: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 4.5.6.7, Local: 1.2.3.4
```

Chapter 18
Fast Ethernet Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Fast Ethernet interfaces.

show interfaces (for Fast Ethernet Interfaces)

Syntax

show interfaces fe-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces fe-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces fe-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about Fast Ethernet router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

fe-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Fast Ethernet interfaces) on page 342
show interfaces brief (for Fast Ethernet interfaces) on page 342
show interfaces detail (for Fast Ethernet interfaces) on page 343
show interfaces extensive (for Fast Ethernet interfaces) on page 344
show interfaces media (for Fast Ethernet interfaces) on page 345
show interfaces statistics (for Fast Ethernet interfaces) on page 346

Output Fields at a Glance

Table 35 summarizes which information is included in the output fields of each show interfaces command option for Fast Ethernet interfaces. In this table, output fields are listed in alphabetical order. In Table 36, the output fields are listed in more detail in the order in which they are displayed.

Table 35: Fast Ethernet Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
Standard Detail Extensive	Active alarms and Active defects—Ethernet-specific defects that can render the interface unable to pass packets.
Extensive	Autonegotiation information—Information about link autonegotiation.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	Current address—Configured MAC address.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Device Flags—Information about the physical device.
Extensive	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.
All	Enabled—State of the interface.
Standard Detail Extensive	Encapsulation—Encapsulation being used on the physical interface.
Extensive	Filter Statistics—Statistics reported by the PIC's filter subsystem.
All	Flow control—Whether flow control is enabled.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	Hardware address—Hardware MAC address.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
All	Ingress Rate Limit—Configured ingress limit. It can be from 1 to 1000 Mbps.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags—Information about the interface.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link flags—Information about the link.

Options	Field Description
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
Extensive	MAC Statistics—Statistics reported by the PIC's MAC subsystem.
All	Minimum links needed—(For aggregate Ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
All	MTU—MTU size on the physical interface.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface and Index—Name of the physical interface and the physical interface's index number, which reflects its initialization sequence.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	Source filtering—Whether source filter is configured.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number of packets and bytes received and transmitted on the physical interface.
Logical Interface	
All	Addresses—Addresses associated with the logical interface.
All	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
All	Family Flags—Information about the protocol family flags.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Local—IP address of the logical interface.
All	Logical interface and Index—Name of the logical interface and the logical interface's index number, which reflects its initialization sequence.
All	Logical Interface Flags—Information about the logical interface.
All	MTU—MTU size on the logical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Statistics—(For aggregate Ethernet only) Information about the number of packets, packets per second, number of bytes and bytes per second on this aggregate interface.

Table 36: Fast Ethernet Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Source filtering	Whether source filter is configured.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (either local or remote).
Flow control	Whether flow control is enabled.
Ingress Rate Limit	Performs port-based rate-limiting to the ingress traffic at the PIC. This value may vary from 1 to 100 Mbps. This feature is only supported for 12- and 48-port Fast Ethernet PICs.
Minimum links needed	(For aggregate Ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Current address	Configured MAC address.
Hardware address	(Extensive output only) Hardware MAC address.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.

Output Field	Output Field Description
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. ■ FIFO errors—Number of FIFO errors in the receive direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably broken.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. ■ FIFO errors—Number of FIFO errors in the receive direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably broken.
Active alarms and Active defects	Ethernet-specific defects 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. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or the turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link. Link means that the 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 broken.
MAC statistics	(Extensive output only) Statistics reported by the PIC's MAC subsystem.

Output Field	Output Field Description
Filter Statistics	(Extensive output only) Statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. ■ Input packet count—Number of packets received from the MAC hardware that the filter processed. ■ Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. ■ Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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 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.
Autonegotiation information	(Extensive output only) Information about link autonegotiation.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Statistics	(For aggregate Ethernet only) Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface. ■ Bundle—Information about input and output bundles. ■ Link—Information about specific links in the aggregate. ■ Marker Statistics—Information about 802.3ad marker protocol statistics on the specified links. ■ Marker Rx—Marker PDUs received. The number of valid marker PDUs received on this aggregation port. ■ Resp Tx—Marker response PDUs transmitted. The number of marker response PDUs transmitted on this aggregation port. ■ Unknown Rx—Unknown received. The number of frames received that either carry the slow protocols Ethernet type value (43B.4), but contain an unknown PDU, or are addressed to the slow protocols group MAC address (43B.3), but do not carry the slow protocols Ethernet type. ■ Illegal Rx—Illegal received. The number of frames received that carry the slow protocols Ethernet type value (43B.4), but contain a badly formed PDU or an illegal value of protocol subtype (43B.4). ■ Addresses—Addresses associated with the logical interface.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
MTU	MTU size on the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for Fast Ethernet interfaces)

```
user@host> show interfaces fe-0/3/3
Physical interface: fe-0/3/3, Enabled, Physical link is Up
  Interface index: 66, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0d:24:60, Hardware address: 00:90:69:0d:24:60
  Last flapped   : 2002-04-17 18:54:15 PDT (00:03:35 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface fe-0/3/3.0 (Index 70) (SNMP ifIndex 112)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.54.33/24, Local: 10.54.33.1, Broadcast: 10.54.33.255
```

show interfaces brief (for Fast Ethernet interfaces)

```
user@host> show interfaces brief fe-0/3/0
Physical interface: fe-0/3/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fe-0/3/0.0
    Flags: SNMP-Traps VLAN 200 Encapsulation: ENET2
    aenet
```

show interfaces detail (for Fast Ethernet interfaces)

```
user@host> show interfaces detail fe-0/3/0
Physical interface: fe-0/3/0, Enabled, Physical link is Up
  Interface index: 63, SNMP ifIndex: 19, Generation: 62
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:0d:27:f2, Hardware address: 00:90:69:0d:24:5d
  Last flapped   : 2002-04-17 18:54:15 PDT (00:05:47 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  84                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   3                    0 pps
   Output packets:                   0                    0 pps
  Active alarms  : None
  Active defects : None

  Logical interface fe-0/3/0.0 (Index 67) (SNMP ifIndex 60) (Generation 66)
    Flags: SNMP-Traps VLAN 200 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  84
     Output bytes  :                   0
     Input  packets:                   3
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                  84                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   3                    0 pps
     Output packets:                   0                    0 pps
    Protocol aenet, AE bundle: ae2.0, MTU: 0, Flags: Generation: 159
    Route table: 0
```

show interfaces extensive (for Fast Ethernet interfaces)

```
user@host> show interfaces extensive fe-0/3/0
Physical interface: fe-0/3/0, Enabled, Physical link is Up
  Interface index: 63, SNMP ifIndex: 19, Generation: 62
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:0d:27:f2, Hardware address: 00:90:69:0d:24:5d
  Last flapped   : 2002-04-17 18:54:15 PDT (00:05:47 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                  84                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   3                    0 pps
   Output packets:                   0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 13, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                          1272                0
    Total packets                           19                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                       19                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                      16
    Input packet rejects                     0
    Input DA rejects                        13
    Input SA rejects                         0
    Output packet count                                      0
    Output packet pad count                                  0
    Output packet error count                                0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: Ok,
    Link partner: Full-duplex, Flow control: None
  Packet Forwarding Engine configuration:
    Destination slot: 0
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %          bps   %        bytes
    0 best-effort             0            0   0            0      low    none
    1 expedited-forwarding    0            0   0            0      low    none
    2 assured-forwarding      0            0   0            0      low    none
    3 network-control         0            0   0            0      low    none
```

```
  Logical interface fe-0/3/0.0 (Index 67) (SNMP ifIndex 60) (Generation 66)
    Flags: SNMP-Traps VLAN 200 Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  84
     Output bytes  :                   0
     Input  packets:                   3
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                  84                    0 bps
     Output bytes  :                   0                    0 bps
     Input  packets:                   3                    0 pps
     Output packets:                   0                    0 pps
    Protocol aenet, AE bundle: ae2.0, MTU: 0, Flags: Generation: 159
    Route table: 0
```

show interfaces media (for Fast Ethernet interfaces)

```
user@host> show interfaces media fe-0/3/0
Physical interface: fe-0/3/0, Enabled, Physical link is Up
  Interface index: 63, SNMP ifIndex: 19
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0d:27:f2, Hardware address: 00:90:69:0d:24:5d
  Last flapped   : 2002-04-17 18:54:15 PDT (00:05:47 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  MAC statistics:
    Input bytes: 1272, Input packets: 19, Output bytes: 0, Output packets: 0
  Filter statistics:
    Filtered packets: 0, Padded packets: 0, Output packet errors: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: Ok,
    Link partner: Full-duplex, Flow control: None
```

show interfaces statistics (for Fast Ethernet interfaces)

```
user@host> show interfaces statistics fe-0/3/0
Physical interface: fe-0/3/0, Enabled, Physical link is Up
  Interface index: 63, SNMP ifIndex: 19
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0d:27:f2, Hardware address: 00:90:69:0d:24:5d
  Last flapped   : 2002-04-17 18:54:15 PDT (00:05:47 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 13, Output errors: 0
  Active alarms  : None
  Active defects : None

  Logical interface fe-0/3/0.0 (Index 67) (SNMP ifIndex 60)
    Flags: SNMP-Traps VLAN 200 Encapsulation: ENET2
  Input packets : 3
  Output packets: 0
    Protocol aenet, AE bundle: ae2.0, MTU: 0, Flags:
```

Chapter 19
Gigabit Ethernet Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Gigabit Ethernet interfaces.

show interfaces (for Gigabit Ethernet Interfaces)

Syntax

show interfaces ge-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces ge-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces ge-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about Gigabit Ethernet router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

ge-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Gigabit Ethernet interfaces) on page 354
show interfaces (standard) (for Aggregated Gigabit Ethernet interfaces) on page 354
show interfaces brief (for Gigabit Ethernet interfaces) on page 355
show interfaces detail (for Gigabit Ethernet interfaces) on page 355
show interfaces detail (for Aggregated Gigabit Ethernet interfaces) on page 356
show interfaces extensive (for Gigabit Ethernet interfaces) on page 357
show interfaces extensive (for Aggregated Gigabit Ethernet interfaces) on page 358
show interfaces media (for Gigabit Ethernet interfaces) on page 359
show interfaces statistics (for Gigabit Ethernet interfaces) on page 359

Output Fields at a Glance

Table 37 summarizes which information is included in the output fields of each show interfaces command option for Gigabit Ethernet interfaces. In this table, output fields are listed in alphabetical order. In Table 38, the output fields are listed in more detail in the order in which they are displayed.

Table 37: Gigabit Ethernet Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
Standard Detail Extensive	Active alarms and Active defects—Ethernet-specific defects that can render the interface unable to pass packets.
Extensive	Autonegotiation information—Information about link autonegotiation.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	Clocking—Reference clock source. It can be Internal or External.
Standard Detail Extensive	Current address—Configured MAC address.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Device Flags Field—Information about the physical device.
Extensive	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.
All	Enabled Field—State of the interface.
Standard Detail Extensive	Encapsulation—Encapsulation being used on the physical interface.
Standard Detail Extensive	FCS—Frame check sequence on the interface (either 16 or 32).
Extensive	Filter Statistics—Statistics reported by the PIC's filter subsystem.
All	Flow control—Whether flow control is enabled.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Standard Detail Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	Hardware address—Hardware MAC address.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Extensive	Input errors—Input errors on the interface.

Options	Field Description
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link flags—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
Extensive	MAC Statistics—Statistics reported by the PIC's MAC subsystem.
All	Minimum links needed—(For aggregate ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
All	MTU—MTU size on the physical interface.
Extensive	Output errors—Output errors on the interface.
Standard Detail Extensive	Payload scrambler—Whether payload scrambling is enabled.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface and Index—Name of the physical interface and the physical interface's index number, which reflects its initialization sequence.
Extensive	Queue counters—Name of queue. It can be 0, 1, 2, or 3.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	Source filtering—Whether source filter is configured.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Logical Interface	
All	Addresses—Addresses associated with the logical interface.
All	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
All	Family Flags Field—Information about the protocol family flags.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Local—IP address of the logical interface.
All	Logical interface and Index—Name of the logical interface and the logical interface's index number, which reflects its initialization sequence.
All	Logical Interface Flags Field—Information about the logical interface.
All	MTU—MTU size on the logical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Statistics—(For aggregate Ethernet only) Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
Detail Extensive	Traffic statistics—Number of packets and bytes received and transmitted on the logical interface.

Table 38: Gigabit Ethernet Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Source filtering	Whether source filter is configured.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (either local or remote).
Flow control	Whether flow control is enabled.
Minimum links needed	(For aggregate Ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Current address	Configured MAC address.
Hardware address	(Extensive output only) Hardware MAC address.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
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).

Output Field	Output Field Description
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. ■ FIFO errors—Number of FIFO errors in the receive direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably broken.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. ■ FIFO errors—Number of FIFO errors in the receive direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably broken.

Output Field	Output Field Description
Queue counters	Name of queue. It can be 0, 1, 2, or 3. ■ Queued packets—Number of queued packets. ■ Transmitted packets—Number of transmitted packets. ■ RED-dropped packets—Number of RED-dropped packets.
Active alarms and Active defects	Ethernet-specific defects 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. 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. Link means that the 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 broken.
MAC statistics	(Extensive output only) Statistics reported by the PIC's MAC subsystem.
Filter Statistics	(Extensive output only) Statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. ■ Input packet count—Number of packets received from the MAC hardware that the filter processed. ■ Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. ■ Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the 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 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.
Autonegotiation information	(Extensive output only) Information about link autonegotiation.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.

Statistics	(For aggregate Ethernet only) Information about the number of packets, packets per second, number of bytes, and bytes per second on this aggregate interface.
	■ Bundle—Information about input and output bundles.
	■ Link—Information about specific links in the aggregate.
	■ Marker Statistics—Information about 802.3ad marker protocol statistics on the specified links.
	  ■ Marker Rx—Marker PDUs received. The number of valid marker PDUs received on this aggregation port.
	  ■ Resp Tx—Marker response PDUs transmitted. The number of marker response PDUs transmitted on this aggregation port.
	  ■ Unknown Rx—Unknown received. The number of frames received that either carry the slow protocols Ethernet type value (43B.4), but contain an unknown PDU; or are addressed to the slow protocols group MAC address (43B.3), but do not carry the slow protocols Ethernet type.
	  ■ Illegal Rx—Illegal received. The number of frames received that carry the slow protocols Ethernet type value (43B.4), but contain a badly formed PDU or an illegal value of protocol subtype (43B.4).
	  ■ Addresses—Addresses associated with the logical interface.

Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
MTU	MTU size on the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.

Output Field	Output Field Description
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for Gigabit Ethernet interfaces)

```
user@host> show interfaces  ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Interface index: 29, SNMP ifIndex: 16
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:0f:c4:f6, Hardware address: 00:90:69:0f:c4:f6
  Last flapped   : 2002-04-17 16:00:43 PDT (02:03:45 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 160 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/0/0.0 (Index 29) (SNMP ifIndex 61)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.51.55/24, Local: 10.51.55.1, Broadcast: 10.51.55.255
```

show interfaces (standard) (for Aggregated Gigabit Ethernet interfaces)

```
user@host> show interfaces ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 13, SNMP ifIndex: 29
  Link-level type: Ethernet, MTU: 4500, Speed: 3000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0f:c7:f1, Hardware address: 00:90:69:0f:c7:f1
  Last flapped   : Never
  Input rate     : 824 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ae1.0 (Index 18) (SNMP ifIndex 73)
    Flags: SNMP-Traps VLAN 100 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :         84949          0       8865941          824
        Output:           948          0        150712            0
    Protocol inet, MTU: 4456, Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.100/24, Local: 10.51.100.1, Broadcast: 10.51.100.255
```

show interfaces brief (for Gigabit Ethernet interfaces)

```
user@host> show interfaces brief ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None

  Logical interface ge-5/0/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.51.55.1/24
```

show interfaces detail (for Gigabit Ethernet interfaces)

```
user@host> show interfaces detail ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Interface index: 29, SNMP ifIndex: 16, Generation: 28
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:0f:c4:f6, Hardware address: 00:90:69:0f:c4:f6
  Last flapped   : 2002-04-17 16:00:43 PDT (02:03:45 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (1d 07:26 ago)
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :              329512                 160 bps
   Input  packets:                   0                   0 pps
   Output packets:                7844                   0 pps
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                7614                 7614                    0
    1 expedited-fo                  0                    0                    0
    2 assured-forw                  0                    0                    0
    3 network-cont                230                  230                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/0/0.0 (Index 29) (SNMP ifIndex 61) (Generation 83)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None, Generation: 106 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.51.55/24, Local: 10.51.55.1, Broadcast: 10.51.55.255,
        Generation: 77
```

show interfaces detail (for Aggregated Gigabit Ethernet interfaces)

```
user@host> show interfaces detail ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 13, SNMP ifIndex: 29, Generation: 12
  Link-level type: Ethernet, MTU: 4500, Speed: 3000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0f:c7:f1, Hardware address: 00:90:69:0f:c7:f1
  Last flapped   : Never
  Statistics last cleared: 2002-04-16 10:37:59 PDT (1d 07:28 ago)
  Traffic statistics:
   Input  bytes  :             13128095                  824 bps
   Output bytes  :              1603149                    0 bps
   Input  packets:               109443                    0 pps
   Output packets:                15163                    0 pps

  Logical interface ae1.0 (Index 18) (SNMP ifIndex 73) (Generation 71)
    Flags: SNMP-Traps VLAN 100 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :         84981          0       8869269          824
        Output:           948          0        150712            0
    Link:
      ge-5/2/0.0
        Input :         84981          0       8869269          824
        Output:             0          0             0            0
      ge-5/2/1.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-5/2/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Marker Statistics:   Marker Rx    Resp Tx    Unknown Rx   Illegal Rx
      ge-5/2/0.0                0          0            0            0
      ge-5/2/1.0                0          0            0            0
      ge-5/2/2.0                0          0   1000000000            0
    Protocol inet, MTU: 4456, Flags: None, Generation: 92 Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.51.100/24, Local: 10.51.100.1, Broadcast: 10.51.100.255,
        Generation: 65
```

show interfaces extensive (for Gigabit Ethernet interfaces)

```
user@host> show interfaces extensive ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Interface index: 29, SNMP ifIndex: 16, Generation: 28
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:0f:c4:f6, Hardware address: 00:90:69:0f:c4:f6
  Last flapped   : 2002-04-17 16:00:43 PDT (02:03:45 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (1d 07:26 ago)
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :              329512                 160 bps
   Input  packets:                   0                   0 pps
   Output packets:                7844                   0 pps
  Input errors:
    Errors: 3776, Drops: 0, Framing errors: 0, Runts: 3776, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0
  Output errors:
    Carrier transitions: 14, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                7614                 7614                    0
    1 expedited-fo                  0                    0                    0
    2 assured-forw                  0                    0                    0
    3 network-cont                230                  230                    0
  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                          8842           502016
    Total packets                         3776             7844
    Unicast packets                          0                0
    Broadcast packets                        0             7598
    Multicast packets                        0              246
    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                       3764
    VLAN tagged frames                       0
    Code violations                        543
  Filter statistics:
    Input packet count                       0
    Input packet rejects                     0
    Input DA rejects                         0
    Input SA rejects                         0
    Output packet count                                   7844
    Output packet pad count                               7844
    Output packet error count                                0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: Ok,
    Link partner: Full-duplex, Flow control: Symmetric/Asymmetric
```

```
  Packet Forwarding Engine configuration:
    Destination slot: 5
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %          bps   %        bytes
    0 best-effort             0            0   0            0      low    none
    1 expedited-forwarding    0            0   0            0      low    none
    2 assured-forwarding      0            0   0            0      low    none
    3 network-control         0            0   0            0      low    none

  Logical interface ge-5/0/0.0 (Index 29) (SNMP ifIndex 61) (Generation 83)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None, Generation: 106 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.51.55/24, Local: 10.51.55.1, Broadcast: 10.51.55.255,
        Generation: 77
```

show interfaces extensive (for Aggregated Gigabit Ethernet interfaces)

```
user@host> show interfaces extensive ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 13, SNMP ifIndex: 29, Generation: 12
  Link-level type: Ethernet, MTU: 4500, Speed: 3000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Current address: 00:90:69:0f:c7:f1, Hardware address: 00:90:69:0f:c7:f1
  Last flapped   : Never
  Statistics last cleared: 2002-04-16 10:37:59 PDT (1d 07:29 ago)
  Traffic statistics:
   Input  bytes  :             13129967                  824 bps
   Output bytes  :              1603149                    0 bps
   Input  packets:               109461                    0 pps
   Output packets:                15163                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0

  Logical interface ae1.0 (Index 18) (SNMP ifIndex 73) (Generation 71)
    Flags: SNMP-Traps VLAN 100 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :         84999          0       8871141          824
        Output:           948          0        150712            0
    Link:
      ge-5/2/0.0
        Input :         84999          0       8871141          824
        Output:             0          0             0            0
      ge-5/2/1.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      ge-5/2/2.0
        Input :             0          0             0            0
        Output:             0          0             0            0
```

```
  Marker Statistics:   Marker Rx    Resp Tx    Unknown Rx   Illegal Rx
    ge-5/2/0.0                0          0            0            0
    ge-5/2/1.0                0          0            0            0
    ge-5/2/2.0                0          0   1000000000            0
  Protocol inet, MTU: 4456, Flags: None, Generation: 92 Route table: 0
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.51.100/24, Local: 10.51.100.1, Broadcast: 10.51.100.255,
      Generation: 65
```

show interfaces media (for Gigabit Ethernet interfaces)

```
user@host> show interfaces media ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Interface index: 29, SNMP ifIndex: 16
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:0f:c4:f6, Hardware address: 00:90:69:0f:c4:f6
  Last flapped   : 2002-04-17 16:00:43 PDT (02:03:45 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 160 bps (0 pps)
  Active alarms  : None
  Active defects : None
  MAC statistics:
    Input bytes: 8842, Input packets: 3776, Output bytes: 502016,
    Output packets: 7844
  Filter statistics:
    Filtered packets: 0, Padded packets: 7844, Output packet errors: 0
  Autonegotiation information:
    Negotiation status: Complete, Link partner status: Ok,
    Link partner: Full-duplex, Flow control: Symmetric/Asymmetric
```

show interfaces statistics (for Gigabit Ethernet interfaces)

```
user@host> show interfaces statistics ge-5/0/0
Physical interface: ge-5/0/0, Enabled, Physical link is Up
  Interface index: 29, SNMP ifIndex: 16
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Current address: 00:90:69:0f:c4:f6, Hardware address: 00:90:69:0f:c4:f6
  Last flapped   : 2002-04-17 16:00:43 PDT (02:03:46 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (1d 07:26 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 160 bps (0 pps)
  Input errors: 3776, Output errors: 0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/0/0.0 (Index 29) (SNMP ifIndex 61)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.51.55/24, Local: 10.51.55.1, Broadcast: 10.51.55.255
```

Chapter 20
GRE Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot GRE interfaces.

show interfaces (for GRE Interfaces)

Syntax

show interfaces gr-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces gr-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces gr-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about GRE router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

gr-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for GRE Interfaces) on page 366
show interfaces brief (for GRE Interfaces) on page 366
show interfaces detail (for GRE Interfaces) on page 367
show interfaces extensive (for GRE Interfaces) on page 368
show interfaces media (for GRE Interfaces) on page 369
show interfaces statistics (for GRE Interfaces) on page 369

Output Fields at a Glance

Table 39 summarizes which information is included in the output fields of each show interfaces command option for GRE interfaces. In this table, output fields are listed in alphabetical order. In Table 40, the output fields are listed in the order in which they are displayed.

Table 39: GRE Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value, value... xx* seconds
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
Extensive	Clocking—Reference clock source. It can be Internal or External.
Extensive	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.
All	Enabled—State of the interface.
All	Flags—Information about the physical device and interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value, value... xx* seconds
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
Extensive	Link flags—Information about the link.
All	Link-level type—Describes the link layer type.
All	MTU—MTU size on the physical in terface.
Extensive	Output errors—Output errors on the interface.
Extensive	Physical info—Information about the physical interface.
All	Physical interface—Name of the physical interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
Standard Detail Extensive	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.

Options	Field Description
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
All	Type—Encapsulation being used on the interface.
Logical Interface	
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Encapsulation—Encapsulation on the logical interface.
All	Flags—Information about the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	IP Header—IP header of the logical interface.
Detail	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Detail	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Table 40: GRE Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Type	Encapsulation being used on the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Flags	Information about the physical device and interface. Possible values are described in "Device Flags Field" on page 190 and in "Interface Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with Radius authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).

Output Field	Output Field Description
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
IP Header	IP header of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Output Field	Output Field Description
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.

show interfaces (standard) (for GRE Interfaces)

```
user@host> show interfaces gr-3/1/0
Physical interface: gr-3/1/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 76
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface gr-3/1/0.0 (Index 10) (SNMP ifIndex 341)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
  Input packets : 21
  Output packets: 22
    Protocol inet, MTU: 4446, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.0.0/30, Local: 10.10.0.1
    Protocol iso, MTU: 4446, Flags: None
```

show interfaces brief (for GRE Interfaces)

```
user@host> show interfaces gr-3/1/0 brief
Physical interface: gr-3/1/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface gr-3/1/0.0
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    inet  10.10.0.1/30
    iso
```

show interfaces detail (for GRE Interfaces)

```
user@host> show interfaces gr-3/1/0 detail
Physical interface: gr-3/1/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 76, Generation: 80
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                94093                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                 1034                    0 pps
   Output packets:                    0                    0 pps

  Logical interface gr-3/1/0.0 (Index 10) (SNMP ifIndex 341) (Generation 25)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    Traffic statistics:
     Input  bytes  :                  591
     Output bytes  :                  978
     Input  packets:                   15
     Output packets:                   16
    Local statistics:
     Input  bytes  :                  591
     Output bytes  :                  978
     Input  packets:                   15
     Output packets:                   16
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4446, Flags: None, Generation: 45 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.0.0/30, Local: 10.10.0.1, Broadcast: Unspecified,
        Generation: 41
    Protocol iso, MTU: 4446, Flags: None, Generation: 46 Route table: 0
```

show interfaces extensive (for GRE Interfaces)

```
user@host> show interfaces gr-3/1/0 extensive
Physical interface: gr-3/1/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 76, Generation: 80
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                93558                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                 1021                    0 pps
   Output packets:                    0                    0 pps

  Logical interface gr-3/1/0.0 (Index 10) (SNMP ifIndex 341) (Generation 25)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:47:df:64:0000000000000000
    Encapsulation: GRE-NULL
    Traffic statistics:
     Input  bytes  :                   92
     Output bytes  :                  244
     Input  packets:                    3
     Output packets:                    4
    Local statistics:
     Input  bytes  :                   92
     Output bytes  :                  244
     Input  packets:                    3
     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
    Protocol inet, MTU: 4446, Flags: None, Generation: 45 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.10.0.0/30, Local: 10.10.0.1, Broadcast: Unspecified,
        Generation: 41
    Protocol iso, MTU: 4446, Flags: None, Generation: 46 Route table: 0
```

show interfaces media (for GRE Interfaces)

```
user@host> show interfaces media gr-0/3/0
Physical interface: gr-0/3/0, Enabled, Physical link is Up
  Interface index: 14, SNMP ifIndex: 35
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Flags: Present Running SNMP-Traps
  Link flags: None, Physical info: IP-over-GRE
  Input rate     : 0 bps (0 pps), Output rate: 0 bps (0 pps)
```

show interfaces statistics (for GRE Interfaces)

```
user@host> show interfaces gr-0/3/0 statistics detail
Physical interface: gr-0/3/0, Enabled, Physical link is Up
  Interface index: 14, SNMP ifIndex: 35
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Clocking: Unspecified
  Flags: Present Running SNMP-Traps
Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0
    Policed discards: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0
  Logical interface gr-0/3/0.0 (Index 10) (SNMP ifIndex 42)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 192.168.4.253:192.168.4.18:47:df:32:00000000
    Encapsulation: GRE-NULL
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                   0 bps
     Output bytes  :                   0                   0 bps
     Input  packets:                   0                   0 pps
     Output packets:                   0                   0 pps

  Logical interface gr-0/3/0.1 (Index 11) (SNMP ifIndex 43)
...
```

Chapter 21
IP-IP Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot IP-IP interfaces.

show interfaces (for IP-IP Interfaces)

Syntax

show interfaces ip-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces ip-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces ip-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about IP-IP router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

ip-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for IP-IP interfaces) on page 376
show interfaces brief (for IP-IP interfaces) on page 377
show interfaces detail (for IP-IP interfaces) on page 377
show interfaces extensive (for IP-IP interfaces) on page 378
show interfaces media (for IP-IP interfaces) on page 378
show interfaces statistics (for IP-IP interfaces) on page 379

Output Fields at a Glance

Table 41 summarizes which information is included in the output fields of each show interfaces command option for IP-IP interfaces. In this table, output fields are listed in alphabetical order. In Table 42, the output fields are listed in more detail in the order in which they are displayed.

Table 41: IP-IP Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value, value... xx* seconds
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
Extensive	Clocking—Reference clock source. It can be Internal or External.
Extensive	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.
All	Enabled—State of the interface.
All	Flags—Information about the physical device and interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value, value... xx* seconds
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
Standard Detail Extensive	Interface index—Physical interface's index number, which reflects its initialization sequence.
All	Link-level type—Describes the link layer type.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	MTU—MTU size on the physical interface.
Extensive	Output errors—Output errors on the interface.
All	Physical interface—Name of the physical interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
Standard Detail Extensive	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.

Options	Field Description
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
All	Type—Encapsulation being used on the interface.
Logical Interface	
All	Encapsulation—Encapsulation on the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Flags—Information about the logical interface.
Standard	Input packets, Output packets—Number of packets received and transmitted on the logical interface.
All	IP Header—IP header of the logical interface.
Detail Extensive	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics.
Detail Extensive	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Table 42: IP-IP Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Type	Encapsulation being used on the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).

Output Field	Output Field Description
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
IP Header	IP header of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Input Packets, Output packets	(Standard output only) Number of packets received and transmitted on the logical interface.

Output Field	Output Field Description
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in “Family Flags Field” on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in “Address Flags Field” on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for IP-IP interfaces)

```
user@host> show interfaces ip-3/1/0
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Interface index: 82, SNMP ifIndex: 77
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-3/1/0.0 (Index 11) (SNMP ifIndex 353)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:4:df:64:00000000
    Encapsulation: IPv4-NULL
  Input packets : 35
  Output packets: 36
    Protocol inet, MTU: 4450, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.0.0/30, Local: 10.11.0.1
```

show interfaces brief (for IP-IP interfaces)

```
user@host> show interfaces ip-3/1/0 brief
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface ip-3/1/0.0
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:4:df:64:00000000
    Encapsulation: IPv4-NULL
    inet  10.11.0.1/30
```

show interfaces detail (for IP-IP interfaces)

```
user@host> show interfaces ip-3/1/0 detail
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Interface index: 82, SNMP ifIndex: 77, Generation: 81
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 1152                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                   30                    0 pps
   Output packets:                    0                    0 pps

  Logical interface ip-3/1/0.0 (Index 11) (SNMP ifIndex 353) (Generation 26)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:4:df:64:00000000
    Encapsulation: IPv4-NULL
    Traffic statistics:
     Input  bytes  :                 1152
     Output bytes  :                 1748
     Input  packets:                   30
     Output packets:                   30
    Local statistics:
     Input  bytes  :                 1152
     Output bytes  :                 1748
     Input  packets:                   30
     Output packets:                   30
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4450, Flags: None, Generation: 47 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.0.0/30, Local: 10.11.0.1, Broadcast: Unspecified,
        Generation: 43
```

show interfaces extensive (for IP-IP interfaces)

```
user@host> show interfaces ip-3/1/0 extensive
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Interface index: 82, SNMP ifIndex: 77, Generation: 81
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Hold-times     : Up 0 ms, Down 0 ms
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 1044                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                   27                    0 pps
   Output packets:                    0                    0 pps

  Logical interface ip-3/1/0.0 (Index 11) (SNMP ifIndex 353) (Generation 26)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:4:df:64:00000000
    Encapsulation: IPv4-NULL
    Traffic statistics:
     Input  bytes  :                 1044
     Output bytes  :                 1636
     Input  packets:                   27
     Output packets:                   28
    Local statistics:
     Input  bytes  :                 1044
     Output bytes  :                 1636
     Input  packets:                   27
     Output packets:                   28
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :                    0                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:                    0                    0 pps
    Protocol inet, MTU: 4450, Flags: None, Generation: 47 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.0.0/30, Local: 10.11.0.1, Broadcast: Unspecified,
        Generation: 43
```

show interfaces media (for IP-IP interfaces)

```
user@host> show interfaces media ip-3/1/0
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Interface index: 82, SNMP ifIndex: 77
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 136 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces statistics (for IP-IP interfaces)

```
user@host> show interfaces statistics ip-3/1/0
Physical interface: ip-3/1/0, Enabled, Physical link is Up
  Interface index: 82, SNMP ifIndex: 77
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-3/1/0.0 (Index 11) (SNMP ifIndex 353)
    Flags: Point-To-Point SNMP-Traps
    IP-Header 10.255.14.114:10.255.14.132:4:df:64:00000000
    Encapsulation: IPv4-NULL
  Input packets : 692
  Output packets: 729
    Protocol inet, MTU: 4450, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.0.0/30, Local: 10.11.0.1
```

Chapter 22

Multilink Services Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Multilink Services interfaces.

show interfaces (for Multilink Services Interfaces)

Syntax

show interfaces ml-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces ml-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces ml-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about Multilink Services router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display extensive interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

ml-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for Multilink Services Interfaces) on page 387
show interfaces brief (for Multilink Services Interfaces) on page 388
show interfaces detail (for Multilink Services Interfaces) on page 388
show interfaces extensive (for Multilink Services Interfaces) on page 389
show interfaces media (for Multilink Services Interfaces) on page 391
show interfaces statistics (for Multilink Services Interfaces) on page 391

Output Fields at a Glance

Table 43 summarizes which information is included in the output fields of each show interfaces command option for Multilink Services interfaces. In this table, output fields are listed in alphabetical order. In the output fields section, they are listed in more detail in the order in which they are displayed.

Table 43: Multilink Services Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Detail Extensive	Assembly exceptions—Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Detail Extensive	Buffering exceptions—Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface.
All	Device flags—Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Extensive	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.
All	Enabled—State of the interface. Possible values are described in "Enabled Field" on page 190.
Detail Extensive	Frame exceptions—Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Detail Extensive	Hardware errors—Information about hardware errors.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface flags—Information about the interface. Possible values are described in "Interface Flags Field" on page 191.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
All	Link-level type—Encapsulation being used on the physical interface. It can be Multilink.
All	MTU—MTU size on the physical interface.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.

Options	Field Description
Logical Interface	
All	Addresses—Addresses associated with the logical interface.
Detail Extensive	Broadcast—Broadcast address on the logical interface.
Detail Extensive	Bundle events—Information about bundle events.
Detail Extensive	Bundle options—Information about configured bundle options.
All	Destination—IP address of the remote side of the connection.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Flags—Information about the one of the following: ■ The logical interface. Possible values are described in "Logical Interface Flags Field" on page 192. ■ The protocol family flags. Possible values are described in "Family Flags Field" on page 192. ■ The address flags. Possible values are described in "Address Flags Field" on page 192.
Detail Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Detail Extensive	Link—Information about links used in the multilink operation.
All	Local—IP address of the logical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	MTU—MTU size on the logical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Statistics—Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.

Output Fields

Physical interface—Name of the physical interface.

Enabled—State of the interface. Possible values are described in "Enabled Field" on page 190.

Physical link—State of the physical interface. It can be Up or Down.

Interface index—Physical interface's index number, which reflects its initialization sequence.

SNMP ifIndex—SNMP index number for the interface.

Generation—A unique number for use by Juniper Networks Customer Support only.

Link-level type—Encapsulation being used on the physical interface. It can be Multilink.

MTU—MTU size on the physical interface.

Speed—Speed at which the interface is running. It can be up to 450 Mbps.

Device flags—Information about the physical device. Possible values are described in "Device Flags Field" on page 190.

Interface flags—Information about the interface. Possible values are described in "Interface Flags Field" on page 191.

Statistics last cleared—Time when the statistics for the interface were last zeroed.

Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.

Frame exceptions—Information about framing exceptions. Includes events recorded under Exception Events for each logical interface.

- Oversized frames—Number of frames received that exceed maximum frame length. Maximum length is 4500 Kb.
- Errored input frames—Number of input frame errors.
- Input on disabled link/bundle—Number of frames received on disabled links. These can result from either an inconsistent configuration, or from a bundle or link being brought up or down with traffic actively flowing through it.
- Output for disabled link/bundle—Number of frames sent for a disabled or unavailable link. These can result from either an inconsistent configuration, or from a bundle being brought up or down with traffic actively flowing through it.

Buffering exceptions—Information about buffering exceptions. Includes events recorded under Exception Events for each logical interface.

- Packet data buffer overflow—Indicates packet buffer memory full. This can occur when the aggregate data rate exceeds PIC capacity.
- Fragment data buffer overflow—Indicates fragment buffer memory full. This can occur when excessive differential delay is experienced across the links within a single bundle, or when the aggregate data-rate exceeds PIC capacity. Check the logical interface exception event counters to determine which bundle is responsible.

Assembly exceptions—Information about assembly exceptions. Includes events recorded under Exception Events for each logical interface.

> **Note**
> An assembly exception does not necessarily indicate an operational problem with the Multilink PIC itself. If multilink-encapsulated traffic is dropped or reordered after a sequence number has been assigned, the assembling multilink interface notes this and records one or more exception events as a result. The multilink interface could drop multilink-encapsulated fragments itself, as a result. Any multilink packets or fragments dropped by the Multilink PIC itself result in packet or fragment drop counts on individual logical interface(s). If the logical interface drop counts are zero, but exception events are seen, the most likely cause is a problem with the individual link interfaces. Even if the logical interface fragment drop counts are nonzero, it is possible that excess differential delay or traffic losses on individual interfaces are the root cause.

- Fragment timeout—Drop-timer expired while a fragment sequence number was outstanding. Occurs only if drop-timer is enabled. This can occur if the differential delay across the links in a bundle exceeds the drop-timer setting, or if a multilink packet is lost in transit while the drop timer is enabled. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible.

- Missing sequence number—A gap was detected in the sequence numbers of fragments on a bundle. These events do not necessarily indicate any problem with the operation of the Multilink PIC itself. If one or more individual links drop traffic, these events can occur. Check the logical interface exception event counters to determine which bundle is responsible.

- Out-of-order sequence number—Two frames with out-of-order sequence numbers within a single link. This event indicates that an individual link within a bundle reordered traffic, causing the multilink interface to be unable to correctly process the resulting stream. Check the logical interface exception event counters to determine which bundle is responsible.

- Out-of-range sequence number—Received frame with out-of-range sequence number. These events can occur when a large amount of multilink-encapsulated traffic is lost, or the multilink peer is reset, so that a large jump in sequence numbers results. A small number of these events may occur when the far-end of a bundle is taken down or brought up. Check the logical interface exception event counters to determine which bundle is responsible.

Hardware errors—Information about hardware errors.

- Data memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks Customer Support.

- Control memory error—A memory error was detected on the interface DRAM. Indicates possible hardware failure. Contact Juniper Networks Customer Support.

Input errors—(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters:

- Errors—Sum of the incoming 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.

- Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.

- Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.

- Giants—Frames received that are larger than the giant threshold.

- Runts—Frames received that are smaller than the runt threshold.

Logical interface, index, SNMP ifIndex—Name of the logical interface, the logical interface index number (which reflects its initialization sequence), and the logical interface SNMP interface index number.

Flags—Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.

Bundle options—Information about configured bundle options.

- MRRU—Configured size of the MRRU (maximum received reconstructed unit). It can be 1500 though 4500 bytes.
- Drop timer period—Configured drop timeout period. It can be 0 through 127 ms. A value of 0 disables the timer. The default setting is 0.
- Sequence number format—Configured size of the sequence header. It can be 12 bits or 24 bits. The default is 24 bits.
- Fragmentation threshold—Configured fragmentation threshold. A value of 0 results in no fragmentation. Non-zero values can be 128 through 16320 bytes, in integer multiples of 64 bytes. The default setting is 0.

Bundle errors—Information about bundle errors.

- Packet drops—Number and byte count of output packets dropped, rather than being encapsulated and sent out of the router as fragments.
- Fragment drops—Number and byte count of input fragments dropped, rather than being reassembled and handled by the router as packets.
- MRRU exceeded—Number of reassembled packets exceeding the MRRU.
- Exception events—Number of exceptional events encountered while handling traffic on the bundle, other than MRRU exceeded errors. These events are categorized under the physical interface Frame exceptions, Buffering exceptions, and Fragment exceptions. Note that the occurrence of exception events does not necessarily indicate that the multilink interface is not operating properly. It is possible for individual link failures to produce exceptional events.

Statistics—Information about fragments and packets received and sent by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets; output packets are segmented into output fragments for transmission out of the router.

- Bundle—Information about bundles.
- Link—Information about links used in the multilink operation.

Protocol—Protocol running on the logical interface.

MTU—MTU size on the logical interface.

Flags—Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.

Generation—Protocol generation. A unique number for use by Juniper Networks Customer Support only.

Addresses—Addresses associated with the logical interface.

Flags—Information about the address flags. Possible values are described in "Address Flags Field" on page 192.

Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.

Policer—Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.

Destination—IP address of the remote side of the connection.

Local—IP address of the logical interface.

Broadcast—Broadcast address on the logical interface.

Generation—Address generation. A unique number for use by Juniper Networks Customer Support only.

show interfaces (standard) (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:58 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             5          0           450            0
        Output:             6          0           499            0
      Packets:
        Input :             5          0           450            0
        Output:            12          0          1202            0
    Protocol inet, MTU: 1500, Flags:
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.2.2, Local: 1.1.2.1
    Protocol iso, MTU: 1500, Flags: Is-Primary
```

show interfaces brief (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0 brief
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

  Logical interface ml-0/3/0.1
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    inet  1.1.2.1          --> 1.1.2.2
    iso
```

show interfaces detail (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0 detail
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196, Generation: 535
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 3535                    0 bps
   Output bytes  :                 4135                    0 bps
   Input  packets:                   87                    0 pps
   Output packets:                  103                    0 pps
  Frame exceptions:
    Oversized frames                 0
    Errored input frames             0
    Input on disabled link/bundle    0
    Output for disabled link/bundle  0
    Queuing drops                    0
  Buffering exceptions:
    Packet data buffer overflow      0
    Fragment data buffer overflow    0
  Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
    Out-of-range sequence number     0
  Hardware errors (sticky):
    Data memory error                0
    Control memory error             0

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
(Generation 402)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Bundle options:
      MRRU                           1524
      Drop timer period              0
      Sequence number format         long (24 bits)
      Fragmentation threshold        0
      Links needed to sustain bundle 1
    Bundle errors:
      Packet drops                   0 (0 bytes)
      Fragment drops                 0 (0 bytes)
      MRRU exceeded                  0
      Exception events               0
```

```
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             5          0           450            0
        Output:             6          0           499            0
      Packets:
        Input :             5          0           450            0
        Output:            12          0          1202            0
    Link:
      t1-0/1/0:11.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:12.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:10.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:14.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:13.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:15.0
        Input :             0          0             0            0
        Output:             1          0            39            0
      t1-0/1/0:8.0
        Input :             0          0             0            0
        Output:             0          0             0            0
      t1-0/1/0:9.0
        Input :             0          0             0            0
        Output:             0          0             0            0
  Destination class              Packets                Bytes
          dcu-class                    0                    0
    Protocol inet, MTU: 1500, Flags: Generation: 752 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.2.2, Local: 1.1.2.1, Broadcast: Unspecified,
        Generation: 1090
    Protocol iso, MTU: 1500, Flags: Is-Primary, Generation: 753 Route
table: 0
```

show interfaces extensive (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0 extensive
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196, Generation: 535
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                 3535                    0 bps
   Output bytes  :                 4135                    0 bps
   Input  packets:                   87                    0 pps
   Output packets:                  103                    0 pps
```

```
  Frame exceptions:
    Oversized frames                 0
    Errored input frames             0
    Input on disabled link/bundle    0
    Output for disabled link/bundle  0
    Queuing drops                    0
  Buffering exceptions:
    Packet data buffer overflow      0
    Fragment data buffer overflow    0
  Assembly exceptions:
    Fragment timeout                 0
    Missing sequence number          0
    Out-of-order sequence number     0
    Out-of-range sequence number     0
  Hardware errors (sticky):
    Data memory error                0
    Control memory error             0

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
(Generation 402)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Bundle options:
      MRRU                           1524
      Drop timer period              0
      Sequence number format         long (24 bits)
      Fragmentation threshold        0
      Links needed to sustain bundle  1
    Bundle errors:
      Packet drops                   0 (0 bytes)
      Fragment drops                 0 (0 bytes)
      MRRU exceeded                  0
      Exception events               0
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             5          0           450            0
        Output:             6          0           499            0
      Packets:
        Input :             5          0           450            0
        Output:            12          0          1202            0
    Link:
      t1-0/1/0:11.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:12.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:10.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:14.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:13.0
        Input :             1          0            90            0
        Output:             1          0            92            0
      t1-0/1/0:15.0
        Input :             0          0             0            0
        Output:             1          0            39            0
      t1-0/1/0:8.0
        Input :             0          0             0            0
        Output:             0          0             0            0
```

```
      t1-0/1/0:9.0
        Input :             0          0             0            0
        Output:             0          0             0            0
  Destination class              Packets                Bytes
          dcu-class                   0                    0
    Protocol inet, MTU: 1500, Flags: Generation: 752 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.2.2, Local: 1.1.2.1, Broadcast: Unspecified,
        Generation: 1090
    Protocol iso, MTU: 1500, Flags: Is-Primary, Generation: 753 Route
table: 0
```

show interfaces media (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0 media
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces statistics (for Multilink Services Interfaces)

```
user@host> show interfaces ml-0/3/0 statistics
Physical interface: ml-0/3/0, Enabled, Physical link is Up
  Interface index: 273, SNMP ifIndex: 196
  Link-level type: Multilink, MTU: 4474
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2002-04-25 14:21:34 PDT (21:06:59 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ml-0/3/0.1 (Index 110) (SNMP ifIndex 674)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
    Bandwidth: 12288kbps
    Statistics         Frames        fps         Bytes          bps
    Bundle:
      Fragments:
        Input :             5          0           450            0
        Output:             6          0           499            0
      Packets:
        Input :             5          0           450            0
        Output:            12          0          1202            0
    Protocol inet, MTU: 1500, Flags:
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 1.1.2.2, Local: 1.1.2.1
    Protocol iso, MTU: 1500, Flags: Is-Primary
```

Chapter 23

SONET/SDH Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot SONET/SDH interfaces.

show interfaces (for SONET/SDH Interfaces)

Syntax

show interfaces so-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces so-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces so-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about SONET/SDH router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

so-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for SDH/SONET interfaces) on page 402
show interfaces brief (for SDH/SONET interfaces) on page 403
show interfaces detail (for SDH/SONET interfaces) on page 403
show interfaces detail (for Aggregated SDH/SONET interfaces) on page 404
show interfaces extensive (for SDH/SONET interfaces) on page 405
show interfaces extensive (for Aggregated SDH/SONET interfaces) on page 407
show interfaces media (for SDH/SONET interfaces) on page 408
show interfaces statistics (for SDH/SONET interfaces) on page 408

Output Fields at a Glance

Table 44 summarizes which information is included in the output fields of each show interfaces command option for SONET/SDH interfaces. In this table, output fields are listed in alphabetical order. In Table 45, the output fields are listed in more detail in the order in which they are displayed.

Table 44: SONET/SDH Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
All	FCS—Frame check sequence on the interface (either 16 or 32).
All	Device Flags Field—Information about the physical device.
Extensive	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.
All	Enabled Field—State of the interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
All	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Detail Extensive	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.

Options	Field Description
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link flags—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	Minimum links needed—(For aggregate Ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
All	MTU—MTU size on the physical interface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
All	Payload scrambler—Whether payload scrambling is enabled.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
Extensive	POS configuration—Information about how the POS interface is configured.
Extensive	Queue counters—Name of queue. It can be 0, 1, 2, or 3.
Extensive	Received path trace, Transmitted path trace—SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link.
Extensive	Received SONET overhead, Transmitted SONET overhead—Values in the received and transmitted SONET/SDH overhead.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
All	SONET alarms and SONET defects—SONET/SDH media-specific defects 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.
Extensive	SONET PHY, SONET section, SONET line, SONET path—Counts of specific SONET/SDH errors.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number of packets and bytes received and transmitted on the physical interface.
Logical Interface	
All	Addresses flags—Information about the address flags.
All	Addresses—Addresses associated with the logical interface.
Detail Extensive	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	DLCI—If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
All	Encapsulation—Encapsulation on the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Local—IP address of the logical interface.
All	Logical interface flags—Information about the logical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.

Options	Field Description
All	MTU—MTU size on the logical interface.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
Detail Extensive	Protocol Family Flags—Information about the protocol family flags.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.

Table 45: SONET/SDH Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	SONET/SDH reference clock source. It can be internal or external. Clocking is configured and displayed only for channel 0.
Speed	Speed at which the interface is running.
Minimum links needed	(For aggregate Ethernet only) Number of child links that must be operational for the aggregate interface to be operational.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Payload scrambler	Whether payload scrambling is enabled.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.

Output Field	Output Field Description
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Keepalive statistics	Information about keepalive packets.
	■ Input—Number of keepalive packets received by PPP.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*.
	■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.

LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Count of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Queue counters	Name of queue. It can be 0, 1, 2, or 3. ■ Queued packets—Number of queued packets. ■ Transmitted packets—Number of transmitted packets. ■ RED-dropped packets—Number of RED-dropped packets.

Output Field	Output Field Description

SONET alarms and SONET defects	SONET/SDH media-specific defects 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. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. The following lists all possible alarms and defects. For complete explanations of most of these alarms and defects, see Chapter 6 in GR-253, *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*.
	■ AIS—Alarm indication signal
	  ■ AIS-L—Alarm indication signal (line)
	  ■ AIS-P—Alarm indication signal (path)
	■ BERR—Bit error rate
	  ■ BERR-SD—Bit error rate defect—signal degrade
	  ■ BERR-SF—Bit error rate fault—signal fail
	■ LOF—Loss of frame
	■ LOL—Loss of light
	■ LOP-P—Loss of pointer (path)
	■ LOS—Loss of signal
	■ PLL—Phase locked loop
	■ PLM-P—Payload label mismatch
	■ RDI—Remote defect indication
	  ■ RDI-L—Remote defect indication (line)
	  ■ RDI-P—Remote defect indication (path)
	■ UNEQ-P—Unequipped (path)

Output Field	Output Field Description

SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Active alarms and defects, plus counts of specific SONET/SDH errors with detailed information.
	■ Seconds—Number of seconds the defect has been active.
	■ Count—Number of times that the defect has gone from inactive to active.
	■ State—State of the error. State other than OK indicates a problem.
	In addition to those listed under active alarms and defects, the SONET PHY, SONET section, SONET line, and SONET path fields report the following values:
	■ BIP—Bit Interleaved Parity
	  ■ BIP-B1—Bit Interleaved Parity for SONET/SDH section overhead
	  ■ BIP-B2—Bit Interleaved Parity for SONET/SDH line overhead
	  ■ BIP-B3—Bit Interleaved Parity for SONET/SDH path overhead
	■ ES—Errored Seconds
	  ■ ES-L—Errored Seconds (near-end line)
	  ■ ES-LFE—Errored Seconds (far-end line)
	  ■ ES-P—Errored Seconds (near-end STS path)
	  ■ ES-PFE—Errored Seconds (far-end STS path)
	  ■ ES-S—Errored Seconds (section)
	■ PHY Light—Loss of optical signal
	■ REI—Remote Error Indication
	  ■ REI-L—Remote Error Indication (line)
	  ■ REI-P—Remote Error Indication (path)
	■ SEF—Severely Errored Frame
	■ SEFS-S—Severely Errored Framing Seconds (section)
	■ SES—Severely Errored Seconds
	  ■ SES-L—Severely Errored Seconds (near-end line)
	  ■ SES-LFE—Severely Errored Seconds (far-end line)
	  ■ SES-P—Severely Errored Seconds (near-end STS path)
	  ■ SES-PFE—Severely Errored Seconds (far-end STS path)
	  ■ SES-S—Severely Errored Seconds (section)
	■ UAS—Unavailable Seconds
	  ■ UAS-L—Unavailable Seconds (near-end line)
	  ■ AS-LFE—Unavailable Seconds (far-end line)
	  ■ UAS-P—Unavailable Seconds (near-end STS path)
	  ■ UAS-PFE— Unavailable Seconds (far-end STS path)

Received SONET overhead, Transmitted SONET overhead	(Extensive output only) Values in the received and transmitted SONET/SDH overhead.
Received path trace, Transmitted path trace	(Extensive output only) POS interfaces allow path trace bytes to be sent inband across the SONET/SDH link. We 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 host name 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, and the transmitted path trace value is the message that this router transmits.

Output Field	Output Field Description
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
POS configuration	(Extensive output only) Information about how the POS interface is configured.
SONET PHY, SONET section, SONET line, SONET path	(Extensive output only) Counts of specific SONET/SDH errors.
Received SONET overhead, Transmitted SONET overhead	(Extensive output only) Values in the received and transmitted SONET/SDH overhead.
Received path trace, Transmitted path trace	(Extensive output only) Channelized OC-12 interfaces allow path trace bytes to be sent inband across the SONET/SDH link. The received path trace value is the message received from the router at the other end of the fiber, and the transmitted path trace value is the message that this router transmits. This information is specific to each of the twelve Channelized OC-12 interfaces.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Output Field	Output Field Description
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for SDH/SONET interfaces)

```
user@host> show interfaces so-0/1/0
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Interface index: 38, SNMP ifIndex: 25
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 10 (00:00:08 ago), Output: 10 (00:00:04 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-04-16 10:17:55 PDT (2d 00:32 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/1/0.0 (Index 30) (SNMP ifIndex 95)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.9.0/30, Local: 10.53.9.1
```

show interfaces brief (for SDH/SONET interfaces)

```
user@host> show interfaces so-0/1/0 brief
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 10 (00:00:08 ago), Output: 10 (00:00:04 ago)
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/1/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  10.53.9.1/30
```

show interfaces detail (for SDH/SONET interfaces)

```
user@host> show interfaces so-0/1/0 detail
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Interface index: 38, SNMP ifIndex: 25, Generation: 37
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 10 (last seen 00:00:08 ago)
    Output: 10 (last sent 00:00:04 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-04-16 10:17:55 PDT (2d 00:32 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (2d 00:12 ago)
  Traffic statistics:
   Input  bytes  :                 397                    0 bps
   Output bytes  :                 412                    0 bps
   Input  packets:                  25                    0 pps
   Output packets:                  26                    0 pps
  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                 25                   25                    0
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/1/0.0 (Index 30) (SNMP ifIndex 95) (Generation 84)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 107 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.9.0/30, Local: 10.53.9.1, Broadcast: Unspecified,
        Generation: 79
```

show interfaces detail (for Aggregated SDH/SONET interfaces)

```
user@host> show interfaces as0 detail
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 44, SNMP ifIndex: 48, Generation: 107
  Link-level type: PPP, MTU: 4474, Speed: 19906560kbps, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                90900                    0 bps
   Output bytes  :                92108                    0 bps
   Input  packets:                 1132                    0 pps
   Output packets:                 1204                    0 pps

  Logical interface as0.0 (Index 40) (SNMP ifIndex 51) (Generation 182)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Aggregate
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1004          0         89356            0
        Output:          1004          0         88352            0
    Link:
      so-0/0/0.0
        Input :          1004          0         89356            0
        Output:          1004          0         88352            0
      so-0/2/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 4470
      Flags: None
    Generation: 223 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.1.0/30, Local: 10.53.1.1, Broadcast: Unspecified,
Generation: 137
```

show interfaces extensive (for SDH/SONET interfaces)

```
user@host> show interface so-0/1/0 extensive
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Interface index: 38, SNMP ifIndex: 25, Generation: 37
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  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 : 10 (last seen 00:00:08 ago)
    Output: 10 (last sent 00:00:04 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-04-16 10:17:55 PDT (2d 00:32 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (2d 00:12 ago)
  Traffic statistics:
   Input  bytes  :                  397                    0 bps
   Output bytes  :                  412                    0 bps
   Input  packets:                   25                    0 pps
   Output packets:                   26                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Bucket drops: 0, Policed discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0,
    HS link FIFO overflows: 1
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0,
    HS link FIFO underflows: 0
  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                 25                   25                    0
  SONET alarms   : None
  SONET defects  : None
  SONET PHY:            Seconds        Count  State
    PLL Lock                  0            0  OK
    PHY Light                 0            0  OK
  SONET section:
    BIP-B1                    0            0
    SEF                       0            0  OK
    LOS                       0            0  OK
    LOF                       0            0  OK
    ES-S                      0
    SES-S                     0
    SEFS-S                    0
```

```
SONET line:
  BIP-B2                    0            0
  REI-L                     0            0
  RDI-L                     0            0  OK
  AIS-L                     0            0  OK
  BERR-SF                   0            0  OK
  BERR-SD                   0            0  OK
  ES-L                      0
  SES-L                     0
  UAS-L                     0
  ES-LFE                    0
  SES-LFE                   0
  UAS-LFE                   0
SONET path:
  BIP-B3                    0            0
  REI-P                     0            0
  LOP-P                     0            0  OK
  AIS-P                     0            0  OK
  RDI-P                     0            0  OK
  UNEQ-P                    0            0  OK
  PLM-P                     0            0  OK
  ES-P                      0
  SES-P                     0
  UAS-P                     0
  ES-PFE                    0
  SES-PFE                   0
  UAS-PFE                   0
Received SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, C2(cmp) : 0xcf, F2      : 0x00
  Z3      : 0x00, Z4      : 0x00, S1(cmp) : 0x00, V5      : 0x00
  V5(cmp) : 0x00
Transmitted SONET overhead:
  F1      : 0x00, J0      : 0x01, K1      : 0x00, K2      : 0x00
  S1      : 0x00, C2      : 0xcf, F2      : 0x00, Z3      : 0x00
  Z4      : 0x00, V5      : 0x00
Received path trace: viper so-2/1/0
  76 69 70 65 72 20 73 6f 2d 32 2f 31 2f 30 00 00   viper so-2/1/0..
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 0d 0a   ................
Transmitted path trace: miura so-0/1/0
  6d 69 75 72 61 20 73 6f 2d 30 2f 31 2f 30 00 00   miura so-0/1/0..
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
HDLC configuration:
  Policing bucket: Disabled
  Shaping bucket : Disabled
  Giant threshold: 4484, Runt threshold: 5
Packet Forwarding Engine configuration:
  Destination slot: 0, PLP byte: 1 (0x00)
  CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                            %          bps   %        bytes
  0 best-effort             0            0   0            0      low    none
  1 expedited-forwarding    0            0   0            0      low    none
  2 assured-forwarding      0            0   0            0      low    none
  3 network-control         0            0   0            0      low    none
```

```
  Logical interface so-0/1/0.0 (Index 30) (SNMP ifIndex 95) (Generation 84)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 107 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.9.0/30, Local: 10.53.9.1, Broadcast: Unspecified,
        Generation: 79
```

show interfaces extensive (for Aggregated SDH/SONET interfaces)

```
user@host> show interfaces as0 extensive
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 44, SNMP ifIndex: 48, Generation: 107
  Link-level type: PPP, MTU: 4474, Speed: 19906560kbps, Minimum links needed: 1
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                91044                    0 bps
   Output bytes  :                92264                   96 bps
   Input  packets:                 1144                    0 pps
   Output packets:                 1216                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0

  Logical interface as0.0 (Index 40) (SNMP ifIndex 51) (Generation 182)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Aggregate
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :          1004          0         89356            0
        Output:          1004          0         88352            0
    Link:
      so-0/0/0.0
        Input :          1004          0         89356            0
        Output:          1004          0         88352            0
      so-0/2/0.0
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet, MTU: 4470
      Flags: None
    Generation: 223 Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.1.0/30, Local:
```

show interfaces media (for SDH/SONET interfaces)

```
user@host> show interfaces so-0/1/0 media
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Interface index: 38, SNMP ifIndex: 25
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 10 (00:00:09 ago), Output: 10 (00:00:05 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-04-16 10:17:55 PDT (2d 00:32 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  SONET alarms   : None
  SONET defects  : None
  SONET errors:
    BIP-B1: 0, BIP-B2: 0, REI-L: 0, BIP-B3: 0, REI-P: 0, BIP-BIP2: 0
  Received path trace: viper so-2/1/0
  Transmitted path trace: miura so-0/1/0
```

show interfaces statistics (for SDH/SONET interfaces)

```
user@host> show interfaces so-0/1/0 statistics
Physical interface: so-0/1/0, Enabled, Physical link is Up
  Interface index: 38, SNMP ifIndex: 25
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 10 (00:00:09 ago), Output: 10 (00:00:05 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-04-16 10:17:55 PDT (2d 00:32 ago)
  Statistics last cleared: 2002-04-16 10:37:59 PDT (2d 00:12 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 1, Output errors: 0
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-0/1/0.0 (Index 30) (SNMP ifIndex 95)
    Flags: Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.53.9.0/30, Local: 10.53.9.1
```

Chapter 24
T1 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot T1 interfaces.

show interfaces (for T1 Interfaces)

Syntax

show interfaces t1-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces t1-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces t1-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about T1 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

t1-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for T1 Interfaces) on page 417
show interfaces brief (for T1 Interfaces) on page 417
show interfaces detail (for T1 Interfaces) on page 417
show interfaces extensive (for T1 Interfaces) on page 418
show interfaces media (for T1 Interfaces) on page 419
show interfaces statistics (for T1 Interfaces) on page 419

Output Fields at a Glance

Table 46 summarizes which information is included in the output fields of each show interfaces command option for T1 interfaces. In this table, output fields are listed in alphabetical order. In Table 47, the output fields are listed in more detail in the order in which they are displayed.

Table 46: T1 Show Interfaces Output Field Summary

Options	Field Description
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—T1 media-specific defects that can render the interface unable to pass packets.
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Buildout—Buildout setting.
All	Clocking—Reference clock source. It can be Internal or External.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Extensive	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.
All	Device Flags Field—Information about the physical device.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Test.
Detail Extensive	DS3 BERT configuration—The BERT (Bit Error Rate Test) checks the quality of the line.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	FCS—Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
All	Framing—Physical layer framing format used on the link. It can be ESF or SF. The default is ESF.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
Detail Extensive	Keepalive settings—Configured settings for keepalives.

Options	Field Description
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	MTU—MTU size on the physical interface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
All	Physical link—State of the physical interface. It can be Up or Down.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Extensive	T1 media—Counts of T1 media-specific errors.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.

Table 47: T1 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Physical link	State of the physical interface. It can be Up or Down.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in "Link Flags Field" on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Output Field	Output Field Description
Keepalive statistics	Information about keepalive packets. ■ Input—Number of keepalive packets received by PPP. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*. ■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received. ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects	T1 media-specific defects 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. 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. ■ LOS—Loss of signal. ■ LOF—Loss of frame. ■ AIS—Alarm indication signal. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.
T1 media	Counts of T1 media-specific errors.

Output Field	Output Field Description
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It can be B8ZS or AMI. ■ Byte encoding—Byte encoding used. It can be Nx64K or Nx56K. ■ Data inversion—HDLC data inversion setting. It can be Enabled or Disabled.
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
Buildout	Buildout setting.
DS1 BERT configuration	BERT (Bit Error Rate Test) checks the quality of the line. This output only appears when BERT is run on the interface (see “test interface bert-start” on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for BERT.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.

show interfaces (standard) (for T1 Interfaces)

```
user@host> show interfaces  t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Interface index: 401, SNMP ifIndex: 3329
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Last flapped   : 2002-04-24 18:37:20 PDT (1d 00:06 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  DS1   alarms   : None
  DS1   defects  : None
```

show interfaces brief (for T1 Interfaces)

```
user@host> show interfaces brief t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
```

show interfaces detail (for T1 Interfaces)

```
user@host> show interfaces detail t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Interface index: 401, SNMP ifIndex: 3329, Generation: 1148
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Last flapped   : 2002-04-24 18:37:20 PDT (1d 00:06 ago)
  Statistics last cleared: 2002-04-24 18:21:06 PDT (1d 00:22 ago)
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                1776                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                 120                    0 pps
  DS1   alarms   : None
  DS1   defects  : None
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: Unknown (0)
```

show interfaces extensive (for T1 Interfaces)

```
user@host> show interfaces extensive t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Interface index: 401, SNMP ifIndex: 3329, Generation: 1148
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Last flapped   : 2002-04-24 18:37:20 PDT (1d 00:06 ago)
  Statistics last cleared: 2002-04-24 18:21:06 PDT (1d 00:22 ago)
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                1776                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                 120                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 1, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, Aged packets: 0
  DS1   alarms   : None
  DS1   defects  : None
  T1  media:            Seconds        Count  State
    SEF                      67            2  OK
    BEE                       0            0  OK
    AIS                       1            1  OK
    LOF                      67            2  OK
    LOS                      66            2  OK
    YELLOW                    0            0  OK
    BPV                       2            2
    EXZ                       2            2
    LCV                       2            6
    PCV                      67        21666
    CS                        0            0
    LES                       2
    ES                       67
    SES                      67
    SEFS                     67
    BES                       2
    UAS                      72
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 1514, Runt threshold: 3
    Timeslots      : All active
    Line encoding: B8ZS, Byte encoding: Nx64K, Data inversion: Disabled
    Buildout       : 0 to 132 feet
  DS1 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Induced Error rate: 10e-0, Algorithm: Unknown (0)
  Packet Forwarding Engine configuration:
    Destination slot: 7, PLP byte: 1 (0x00)
    CoS transmit queue             Bandwidth           Buffer Priority   Limit
                                %          bps   %        bytes
    0 best-effort              95      1459200  95            0      low    none
    3 network-control           5        76800   5            0      low    none
```

show interfaces media (for T1 Interfaces)

```
user@host> show interfaces media t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Interface index: 401, SNMP ifIndex: 3329
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Last flapped   : 2002-04-24 18:37:20 PDT (1d 00:06 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  DS1   alarms   : None
  DS1   defects  : None
  T1 errors:
    BPV: 2, EXZ: 2, LCV: 6, PCV: 21666
    CS: 0
```

show interfaces statistics (for T1 Interfaces)

```
user@host> show interfaces statistics t1-7/3/0
Physical interface: t1-7/3/0, Enabled, Physical link is Up
  Interface index: 401, SNMP ifIndex: 3329
  Link-level type: PPP, MTU: 1504, Clocking: Internal, Speed: T1,
  Loopback: None, FCS: 16, Framing: ESF
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Last flapped   : 2002-04-24 18:37:20 PDT (1d 00:06 ago)
  Statistics last cleared: 2002-04-24 18:21:06 PDT (1d 00:22 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 1, Output errors: 0
  DS1   alarms   : None
  DS1   defects  : None
```

Chapter 25
T3 Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot T3 interfaces.

show interfaces (for T3 Interfaces)

Syntax

show interfaces t3-*fpc/pic/port* <brief | detail | extensive> <*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces t3-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces t3-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about T3 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

t3-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for T3 interfaces) on page 429
show interfaces brief (for T3 interfaces) on page 429
show interfaces detail (for T3 interfaces) on page 430
show interfaces extensive (for T3 interfaces) on page 431
show interfaces media (for T3 interfaces) on page 432
show interfaces statistics (for T3 interfaces) on page 433

Output Fields at a Glance

Table 48 summarizes which information is included in the output fields of each show interfaces command option for T3 interfaces. In this table, output fields are listed in alphabetical order. In Table 49, the output fields are listed in more detail in the order in which they are displayed.

Table 48: T3 Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value*, *value*... *xx* seconds
All	Active alarms and Active defects—T3 media-specific defects that can render the interface unable to pass packets.
Detail Extensive	BERT time period—The configured total time period that the BERT test is to run.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
All	Clocking—Reference clock source. It can be Internal or External.
All	Device Flags Field—Information about the physical device.
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Detail Extensive	DS1 BERT Algorithm—Bit Error Rate Testing.
Detail Extensive	DS3 BERT configuration—The BERT (Bit Error Rate Test) checks the quality of the line.
Extensive	DS3 media—Counts of T3 media-specific errors.
Extensive	DSU configuration—Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface.
Detail Extensive	Elapsed—Actual time elapsed since the start of the BERT.
All	Enabled Field—State of the interface.
All	FCS—Frame check sequence on the interface (either 16 or 32).
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	HDLC configuration—Information about the HDLC configuration.
Detail Extensive	Hold-times—Current interface hold-time up and hold-time down, in milliseconds.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value*, *value*... *xx* seconds
Detail Extensive	Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern.
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, Output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface Flags Field—Information about the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.

Options	Field Description
Extensive	Interface transmit queues—Name of the transmit queues and their associated statistics.
Standard	Keepalive Input, Output—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	Keepalive settings—Configured settings for keepalives.
Detail Extensive	Keepalive statistics—Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
All	Link Flags Field—Information about the link.
All	Link-level type—Encapsulation being used on the physical interface.
All	Loopback—Whether loopback is enabled and the type of loopback (either local or remote).
All	Mode—Whether C-bit parity or M13 mode is enabled.
All	MTU—MTU size on the physical interface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	PFE configuration—Information about how the Packet Forwarding Engine is configured.
All	Physical interface—Name of the physical interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.
All	SNMP ifIndex—SNMP index number for the interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Logical Interface	
All	Address flags—Information about the address.
All	Addresses—Addresses associated with the logical interface.
Detail Extensive	Broadcast—Broadcast address.
All	Destination—IP address of the remote side of the connection.
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	DLCI—If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
All	Encapsulation—Encapsulation on the logical interface.
Standard	Input packets, Output packets—Number of packets received and transmitted on the logical interface.
Detail Extensive	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine.
All	Local—IP address of the logical interface.
All	Logical interface flags—Information about the logical interface.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number, and the logical interface's SNMP interface index number.
Detail Extensive	MTU Family Flags Field—Information about the protocol family flags.

Options	Field Description
All	MTU—MTU size on the logical interface.
All	Protocol—Protocol running on the logical interface.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Detail Extensive	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics.
Detail Extensive	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.

Table 49: T3 Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in “Enabled Field” on page 190.
Interface index	Physical interface’s index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
Mode	Whether C-bit parity mode or M13 mode is enabled.
FCS	Frame check sequence on the interface (either 16 or 32). The default is 16-bit.
Framing	Physical layer framing format used on the link. It can be G704, G704-NO-CRC4, or Unframed. The default is G704.
Device flags	Information about the physical device. Possible values are described in “Device Flags Field” on page 190.
Interface flags	Information about the interface.
Link flags	Information about the link. Possible values are described in “Link Flags Field” on page 191.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value*, *value*... *xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.

Output Field	Output Field Description
Keepalive Input, Output	(Standard output only) Number of keepalive packets sent and received by PPP and how long ago the last keepalive packets were sent and received.
Keepalive settings	Configured settings for keepalives. ■ interval *seconds*—The time in seconds between successive keepalive requests. The range is 10 seconds through 32767 seconds, with a default of 10 seconds. ■ down-count *number*—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is 1 through 255, with a default of 3. ■ up-count *number*—The number of keepalive packets a destination must receive to change a link's status from down to up. The range is 1 through 255, with a default of 1.

Keepalive statistics	Information about keepalive packets.
	■ Input—Number of keepalive packets received by PPP.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was received in the format *hh:mm:ss*.
	■ Output—Number of keepalive packets sent by PPP and how long ago the last keepalive packets were sent and received.
	  ■ (last seen 00:00:00 ago)—Time since the last keepalive packet was sent in the format *hh:mm:ss*.

LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.

Output Field	Output Field Description
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.
Active alarms and Active defects	T3 media-specific defects 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. 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. ■ AIS—Alarm indicator signal. ■ EXZ—Excessive zeros. ■ FERF—Far-end failure. ■ IDLE—Idle alarm. ■ LCV—Line code violation. ■ LOF—Loss of frame. ■ LOS—Loss of signal. ■ PLL—Phase-locked loop out of lock. ■ YLW—Yellow alarm. Indicates errors at the remote site receiver.
DS3 media	(Extensive output only) Counts of T3 media-specific errors.

Output Field	Output Field Description
HDLC configuration	(Extensive output only) Information about the HDLC configuration. ■ Policing bucket—Configured state of the Rx policer. ■ Shaping bucket—Configured state of the Tx shaper. ■ Giant threshold—Giant threshold programmed into the hardware. ■ Runt threshold—Runt threshold programmed into the hardware. ■ Timeslots—Configured time slots for the interface. ■ Line encoding—Line encoding used. It is always HDB3.
Interface transmit queues	Names of the transmit queues and their associated statistics. ■ B/W—Queue bandwidth as a percentage of the total interface bandwidth. ■ WRR—Weighted round robin (in %). ■ Packets—Number of packets transmitted. ■ Drops—Number of packets dropped. ■ Errors—Number of packet errors.
DSU configuration	Information about the DSU configuration. The last three lines (Bit count, Error bit count, and LOS information) are displayed only if a BERT test has ever been run on the interface. ■ Compatibility mode—CSU/DSU compatibility mode. It can be None, Larscom, Kentrox, or Digital-Link. ■ Scrambling—Payload scrambling. It can be Enabled or Disabled. ■ Subrate—Configured subrate setting. Applies only when Digital-Link compatibility mode is used. It can be Disabled or display units in kbps.
Buildout	Buildout setting.
DS3 BERT configuration	BERT (Bit Error Rate Testing) checks the quality of the line. This output only appears when BERT is run on the interface (see “test interface bert-start” on page 459). ■ BERT time period—Configured total time period that the BERT test is to run. ■ Elapsed—Actual time elapsed since the start of the BERT (in seconds). ■ Induced error rate—Configured rate at which the bit errors are induced in the BERT pattern. ■ Algorithm—Type of algorithm selected for the BERT.

Packet Forwarding Engine configuration	(Extensive output only) Information about the configuration of the Packet Forwarding Engine:
	■ Destination slot—FPC slot number.
	■ PLP byte
	■ Stream number—Stream used by the ASIC on the FPC.
	■ CoS transmit queue—The 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 Bytes—Number of bytes allocated to the queue. This value is only nonzero if the buffer size is configured in terms of time.
	  ■ Priority—Queue priority. Possible values are low and high.
	  ■ Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue will only transmit up to the configured bandwidth, even if there is excess bandwidth available. If none is configured, the queue will transmit beyond the configured bandwidth if there is bandwidth available.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
DLCI	If Frame Relay encapsulation is configured, the DLCI number of the logical interface.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For exampl, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for T3 interfaces)

```
user@host> show interfaces  t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 848
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 45 (00:00:03 ago), Output: 45 (00:00:11 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:27 PDT (00:09:12 ago)
  Input rate     : 296 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface t3-1/2/0.0 (Index 9) (SNMP ifIndex 858)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.20.111.0/30, Local: 10.20.111.2
```

show interfaces brief (for T3 interfaces)

```
user@host> show interfaces brief t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 45 (00:00:03 ago), Output: 45 (00:00:11 ago)
  Active alarms  : None
  Active defects : None

  Logical interface t3-1/2/0.0
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    inet  10.20.111.2/30
```

show interfaces detail (for T3 interfaces)

```
user@host> show interfaces detail t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 848, Generation: 80
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  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 : 45 (last seen 00:00:03 ago)
    Output: 45 (last sent 00:00:11 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:27 PDT (00:09:12 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:07:42 ago)
  Traffic statistics:
   Input  bytes  :                 2448                  296 bps
   Output bytes  :                 4415                   48 bps
   Input  packets:                   87                    0 pps
   Output packets:                  135                    0 pps
  Active alarms  : None
  Active defects : None

  Logical interface t3-1/2/0.0 (Index 9) (SNMP ifIndex 858) (Generation 9)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 17 Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.20.111.0/30, Local: 10.20.111.2, Broadcast: Unspecified,
        Generation: 20
```

show interfaces extensive (for T3 interfaces)

```
user@host> show interfaces extensive t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 848, Generation: 80
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  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 : 45 (last seen 00:00:03 ago)
    Output: 45 (last sent 00:00:11 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:27 PDT (00:09:12 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:07:42 ago)
  Traffic statistics:
   Input  bytes  :                 2448                  296 bps
   Output bytes  :                 4415                   48 bps
   Input  packets:                   87                    0 pps
   Output packets:                  135                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 1, Bucket drops: 0,
    Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, HS link CRC errors: 0, SRAM errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0
  Active alarms  : None
  Active defects : None
  DS3 media:            Seconds        Count  State
    PLL Lock                  0            0  OK
    Reframing                 0            0  OK
    AIS                       0            0  OK
    LOF                       0            0  OK
    LOS                       0            0  OK
    IDLE                      0            0  OK
    YELLOW                    0            0  OK
    BPV                       0            0
    EXZ                       0            0
    LCV                       0            0
    PCV                       0            0
    CCV                       0            0
    LES                       0
    PES                       0
    PSES                      0
    CES                       0
    CSES                      0
    SEFS                      0
    UAS                       0
  HDLC configuration:
    Policing bucket: Disabled
    Shaping bucket : Disabled
    Giant threshold: 4484, Runt threshold: 3
  DSU configuration:
    Compatibility mode: Digital Link, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
```

```
  DS-3 BERT configuration:
    BERT time period: 10 seconds, Elapsed: 0 seconds
    Algorithm: Unknown (0), Induced error rate: 10e-0
  Packet Forwarding Engine configuration:
    Destination slot: 1, PLP byte: 1 (0x00)
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                              %          bps   %        bytes
    0 best-effort             0            0   0            0      low    none
    1 expedited-forwarding    0            0   0            0      low    none
    2 assured-forwarding      0            0   0            0      low    none
    3 network-control         0            0   0            0      low    none

  Logical interface t3-1/2/0.0 (Index 9) (SNMP ifIndex 858) (Generation 9)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None, Generation: 17 Route table: 0
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.20.111.0/30, Local: 10.20.111.2, Broadcast: Unspecified,
        Generation: 20
```

show interfaces media (for T3 interfaces)

```
user@host> show interfaces media t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 848
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 45 (00:00:03 ago), Output: 46 (00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:27 PDT (00:09:12 ago)
  Input rate     : 296 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  Active alarms  : None
  Active defects : None
  DS3 errors:
    BPV: 0, EXZ: 0, LCV: 0, PCV: 0
    CCV: 0
```

show interfaces statistics (for T3 interfaces)

```
user@host> show interfaces statistics t3-1/2/0
Physical interface: t3-1/2/0, Enabled, Physical link is Up
  Interface index: 81, SNMP ifIndex: 848
  Link-level type: PPP, MTU: 4474, Clocking: Internal, Speed: T3,
  Loopback: None, FCS: 16, Mode: C/Bit parity
  Device flags   : Present Running
  Interface flags: Link-Layer-Down Point-To-Point SNMP-Traps
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 45 (00:00:03 ago), Output: 46 (00:00:00 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Not-configured
  Last flapped   : 2002-05-16 15:34:27 PDT (00:09:12 ago)
  Statistics last cleared: 2002-05-16 15:35:57 PDT (00:07:42 ago)
  Input rate     : 296 bps (0 pps)
  Output rate    : 48 bps (0 pps)
  Input errors: 1, Output errors: 0
  Active alarms  : None
  Active defects : None

  Logical interface t3-1/2/0.0 (Index 9) (SNMP ifIndex 858)
    Flags: Device-Down Point-To-Point SNMP-Traps Encapsulation: PPP
    Protocol inet, MTU: 4470, Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.20.111.0/30, Local: 10.20.111.2
```

Chapter 26
Virtual Tunnel Interfaces Monitoring and Troubleshooting

This chapter summarizes the command-line interface (CLI) commmands you can use to monitor and troubleshoot Virtual Tunnel (VT) interfaces.

show interfaces (for VT Interfaces)

Syntax

show interfaces t3-*fpc/pic/port* <brief | detail | extensive><*interface-name* destination-class *destination-class-name*> <*interface-name* source-class *source-class-name*>

show interfaces t3-*fpc/pic/port* <brief | detail | extensive> <media>

show interfaces t3-*fpc/pic/port* <brief | detail | extensive> <statistics>

Description

Display status information about T3 router interfaces.

Options

none—Display information about all interfaces.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display very detailed interface information.

destination-class *destination-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the destination address matching those prefixes. Whenever a destination class is specified, you must also specify a particular logical interface, not all interfaces.

t3-*fpc/pic/port*—Name of an interface.

media—(Optional) Display media-specific information about network interfaces.

source-class *source-class-name*—(Optional) Name of a logical grouping of prefixes that count packets having the source address matching those prefixes. Whenever a source class is specified, you must also specify a particular logical interface, not all interfaces.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Sample Output

show interfaces (standard) (for VT interfaces) on page 441
show interfaces brief (for VT interfaces) on page 441
show interfaces detail (for VT interfaces) on page 442
show interfaces extensive (for VT interfaces) on page 442
show interfaces media (for VT interfaces) on page 443
show interfaces statistics (for VT interfaces) on page 443

Output Fields at a Glance

Figure 50 summarizes which information is included in the output fields of each show interfaces command option for VT interfaces. In this table, output fields are listed in alphabetical order. In Table 51, the output fields are listed in more detail in the order in which they are displayed.

Table 50: VT Show Interfaces Output Field Summary

Options	Field Description
Physical Interface	
Extensive	ANSI LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ANSI LMI settings: *value, value... xx* seconds
Extensive	Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off.
Standard Detail Extensive	CHAP state—Displays the state of the challenge-handshake protocol during its transaction.
Extensive	Clocking—Reference clock source. It can be Internal or External.
Extensive	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.
All	Enabled—State of the interface.
All	Flags—Information about the physical device and interface.
Extensive	Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors.
Detail Extensive	Generation—A unique number for use by Juniper Networks Customer Support only.
Extensive	Giants—Frames received that are larger than the giant threshold.
Extensive	ITU LMI settings—Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is ITU LMI settings: *value, value... xx* seconds
Extensive	Input errors—Input errors on the interface.
Standard	Input rate, output rate—Rate of bits and packets received and transmitted on the interface.
All	Interface index—Physical interface's index number, which reflects its initialization sequence.
All	LCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	LMI Statistics—Statistics about link management, including a count of packets sent and received, and the time of the last activity.
All	Last flapped—Date, time, and how long ago the interface went from down to up.
Extensive	Link flags—Information about the link.
All	Link-level type—Describes the link layer type.
All	MTU—MTU size on the physical in terface.
All	NCP state—Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Extensive	Output errors—Output errors on the interface.
Extensive	Physical info—Information about the physical interface.
All	Physical interface—Name of the physical interface.
Extensive	Runts—Frames received that are smaller than the runt threshold.

Options	Field Description
Standard Detail Extensive	SNMP ifIndex—SNMP index number for the interface.
All	Speed—Speed at which the interface is running.
Detail Extensive	Statistics last cleared—Time when the statistics for the interface were last zeroed.
Detail Extensive	Traffic statistics—Number and rate of bytes and packets received and transmitted on the physical interface.
All	Type—Encapsulation being used on the interface.
Logical Interface	
Detail Extensive	Destination class—List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
All	Encapsulation—Encapsulation on the logical interface.
Detail Extensive	Filters—Name of the firewall filters to be evaluated when packets are received or transmitted on the interface.
All	Flags—Information about the logical interface.
All	IP Header—IP header of the logical interface.
Detail	Local statistics—Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
All	Logical interface, Index, SNMP ifIndex—Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Detail Extensive	Policer—Policers to be evaluated when packets are received or transmitted on the interface.
Detail Extensive	RPF Failures: Packets: xx, Bytes: yy—The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Detail	Traffic statistics—Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail	Transit statistics—Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Detail Extensive	Source class—List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.

Table 51: VT Show Interfaces Output Field Summary in Order of Appearance

Output Field	Output Field Description
Physical Interface	
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in "Enabled Field" on page 190.
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Generation	A unique number for use by Juniper Networks Customer Support only.
Type	Encapsulation being used on the interface.
Link-level type	Encapsulation being used on the physical interface.
MTU	MTU size on the physical interface.
Clocking	Reference clock source. It can be Internal or External.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback (local or remote).
Device flags	Information about the physical device. Possible values are described in "Device Flags Field" on page 190.
Interface flags	Information about the interface.
LMI settings	(Extensive output only) Settings for link management can be either ANSI LMI settings or ITU LMI settings. ANSI LMI settings is the default. The format is (ANSI or ITU) LMI settings: *value, value... xx* seconds, where *value* can be: ■ n391dte—DTE full status polling interval (1..255) ■ n392dce—DCE error threshold (1..10) ■ n392dte—DTE error threshold (1..10) ■ n393dce—DCE monitored event count (1..10) ■ n393dte—DTE monitored event count (1..10) ■ t391dte—DTE polling timer (5..30 seconds) ■ t392dce—DCE polling verification timer (5..30 seconds)
LMI Statistics	(Extensive output only) Statistics about the link management. ■ Input—Number of packets coming in on the interface (*nn*) and how much time has passed since the last packet arrived. The format is Input: *nn* (last seen hh:mm:ss ago). ■ Output—Number of packets sent out on the interface (*nn*) and how much time has passed since the last packet was sent. The format is Output: *nn* (last sent *hh:mm:ss* ago).
LCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
NCP state	Specific PPP bits. Opened indicates that they have been initialized and opened, which means that the link is healthy.
Statistics last cleared	Time when the statistics for the interface were last zeroed.
CHAP state	Displays the state of the challenge-handshake protocol during its transaction. ■ Not-configured—CHAP was not configured on the interface. ■ Success—CHAP authentication was successful. ■ Fail—CHAP authentication failed. ■ Chap-Resp-received—Received response for the challenge sent, but not yet moved into the Success state. (Most likely with RADIUS authentication.) ■ Chap-Resp-sent—Response sent for the challenge received. ■ Chap-Chal-sent—Challenge sent. ■ Chap-Chal-received—Challenge received but response not yet sent.
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).

Output Field	Output Field Description
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. ■ Input bytes, Output bytes—Number of bytes received and transmitted on the interface. ■ Input packets, Output packets—Number of packets received and transmitted on the interface.
Input rate, Output rate	(Standard output only) Rate of bits (in bbs) and packets (in pps) received and transmitted on the interface.
Input errors	(Extensive output only) Input errors on the interface. The following paragraphs explain the nonobvious counters: ■ Errors—Sum of the incoming 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. ■ Invalid VCs—Number of cells that arrived for a nonexistent VC. ■ Framing errors—Sum of AAL5 packets that have FCS errors, AAL5 packets that have reassembly timeout errors, and AAL5 packets that have length errors. ■ Bucket Drops—Drops due to traffic load exceeding the interface transmit/receive leaky bucket configuration. The default is off. ■ Giants—Frames received that are larger than the giant threshold. ■ Runts—Frames received that are smaller than the runt threshold. ■ Policed discards—Frames that the incoming packet match code discarded because they were not recognized or of interest. Usually, this field reports protocols that the JUNOS software does not handle, such as CDP. ■ L3 incompletes—Increments when the incoming packet fails Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header would be discarded and this counter would increment. ■ L2 channel errors—This counter increments when the software could not find a valid logical interface for an incoming frame. ■ L2 mismatch timeouts—Count of malformed or short packets that cause the incoming packet handler to discard the frame as unreadable. ■ SRAM errors—This counter increments when a hardware error has occurred in the SRAM on the PIC. The value in this field should always be 0. If it increments, the PIC is broken. ■ HS link FCS errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces.
Output errors	(Extensive output only) Output errors on the interface. The following paragraphs explain the nonobvious counters: ■ Carrier transitions—Number of times the interface has gone from down to up. This number should not increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or a similar problem occurs. If it increments quickly (perhaps once every 10 seconds), then either the cable, the far-end system, or the PIC is broken. ■ 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. ■ Aged packets—Number of packets that remained in shared packet SDRAM for so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly broken hardware.

Output Field	Output Field Description
Logical Interface	
Logical interface, Index, SNMP ifIndex	Name of the logical interface, the logical interface's index number (which reflects its initialization sequence), and the logical interface's SNMP interface index number.
Flags	Information about the logical interface. Possible values are described in "Logical Interface Flags Field" on page 192.
Protocol	Protocol running on the logical interface, such as iso, inet6, mpls.
Encapsulation	Encapsulation on the logical interface.
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize. ■ Input rate—Rate of bits and packets received on the interface. ■ Output rate—Rate of bits and packets transmitted on the interface.
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes a while (generally, less than 1 second) for this counter to stabilize.
MTU	MTU size on the logical interface.
Flags	Information about the protocol family flags. Possible values are described in "Family Flags Field" on page 192.
Generation	A unique number for use by Juniper Networks Customer Support only.
Route table	The address is located in this route table. For example, Route table:0 refers to inet.0.
Filters	Name of the firewall filters to be evaluated when packets are received or transmitted on the interface. The format is Filters: Input: *input-filter-name*, Output: *output-filter-name*.
RPF Failures: Packets: xx, Bytes: yy	The amount of incoming traffic (in packets and bytes) that failed a unicast Reverse Path Forwarding (RPF) check on this interface.
Destination class	List of the names of destination class usage (DCU) counters per family and per class for this interface. The counters display Packets and Bytes going to designated user-selected prefixes.
Source class	List of the names of source class usage (SCU) counters per family and per class for this interface. The counters display Packets and Bytes arriving from designated user-selected prefixes.
Policer	Policers to be evaluated when packets are received or transmitted on the interface. The format is Policer: Input: *type-fpc/pic/port*-in-policer, Output: *type-fpc/pic/port*-out-policer.
Addresses	Addresses associated with the logical interface.
Flags	Information about the address flags. Possible values are described in "Address Flags Field" on page 192.
Destination	IP address of the remote side of the connection.
Local	IP address of the logical interface.
Broadcast	Broadcast address.

show interfaces (standard) (for VT interfaces)

```
user@host> show interfaces vt-1/3/0
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Interface index: 36, SNMP ifIndex: 50
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Last flapped   : 2002-04-18 10:40:21 PDT (04:45:03 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface vt-1/3/0.0 (Index 7) (SNMP ifIndex 23)
    Flags: Point-To-Point SNMP-Traps Encapsulation: VPN-loopback-tunnel
    Protocol inet, MTU: Unlimited, Flags: None
    Protocol mpls, MTU: Unlimited, Flags: None
```

show interfaces brief (for VT interfaces)

```
user@host> show interfaces vt-1/3/0 brief
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface vt-1/3/0.0
    Flags: Point-To-Point SNMP-Traps Encapsulation: VPN-loopback-tunnel
    inet
    mpls
```

show interfaces detail (for VT interfaces)

```
user@host> show interfaces vt-1/3/0 detail
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Interface index: 36, SNMP ifIndex: 50, Generation: 35
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2002-04-18 10:40:21 PDT (04:45:07 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps

  Logical interface vt-1/3/0.0 (Index 7) (SNMP ifIndex 23) (Generation 8)
    Flags: Point-To-Point SNMP-Traps Encapsulation: VPN-loopback-tunnel
    Protocol inet, MTU: Unlimited, Flags: None, Generation: 16 Route table: 1
    Protocol mpls, MTU: Unlimited, Flags: None, Generation: 17 Route table: 1
```

show interfaces extensive (for VT interfaces)

```
user@host> show interfaces vt-1/3/0 extensive
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Interface index: 36, SNMP ifIndex: 50, Generation: 35
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2002-04-18 10:40:21 PDT (04:45:09 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed
discards: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0

  Logical interface vt-1/3/0.0 (Index 7) (SNMP ifIndex 23) (Generation 8)
    Flags: Point-To-Point SNMP-Traps Encapsulation: VPN-loopback-tunnel
    Protocol inet, MTU: Unlimited, Flags: None, Generation: 16 Route table: 1
    Protocol mpls, MTU: Unlimited, Flags: None, Generation: 17 Route table: 1
```

show interfaces media (for VT interfaces)

```
user@host> show interfaces vt-1/3/0 media
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Interface index: 36, SNMP ifIndex: 50
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Last flapped   : 2002-04-18 10:40:21 PDT (04:45:24 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
```

show interfaces statistics (for VT interfaces)

```
user@host> show interfaces vt-1/3/0 statistics
Physical interface: vt-1/3/0, Enabled, Physical link is Up
  Interface index: 36, SNMP ifIndex: 50
  Type: Software-Pseudo, Link-level type: VPN-loopback-tunnel, MTU: Unlimited,
Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link flags     : None
  Last flapped   : 2002-04-18 10:40:21 PDT (04:45:27 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0

  Logical interface vt-1/3/0.0 (Index 7) (SNMP ifIndex 23)
    Flags: Point-To-Point SNMP-Traps Encapsulation: VPN-loopback-tunnel
    Protocol inet, MTU: Unlimited, Flags: None
    Protocol mpls, MTU: Unlimited, Flags: None
```

Chapter 27

Real-Time Router Interface Monitoring and Troubleshooting

Table 52 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot router interfaces in real-time. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 52: Commands for Real-Time Interface Monitoring and Troubleshooting

Task or Information to Monitor	CLI Command
Interface statistics.	monitor interface on page 449
tcpdump.	monitor traffic on page 452
Automatic Protection Switching.	show aps on page 455

Monitor Traffic Match Conditions

In the monitor traffic command, you can specify an expression to match by using the matching option and including the expression in quotation marks:

monitor traffic matching *"expression"*

expression is one or more of the match conditions listed in Table 53. To combine expressions, use the logical operators listed in Table 54. For match conditions that have a numeric value, you can specify them using the arithmetic and relational operators listed in Table 55.

Table 53: Match Conditions for the monitor traffic Command

Match Condition	Description
Entity Type	Entity types match packets according to host names and addresses, network addresses, port names and numbers, and subnet masking.
host [*address* | *hostname*]	Matches packets that contain the specified address or hostname. The host match condition can be prepended with the protocol match conditions arp, ip, or rarp, or any of the directional match conditions.
network *address*	Matches packets with source or destination addresses containing the specified network number.
network *address* mask *mask*	Matches packets containing the specified network address and subnet mask.

Match Condition	Description
port [*port-number* | *port-name*]	Matches packets containing the specified source or destination TCP or UDP port number or port name. In place of the numeric port address, you can specify a text synonym (the port numbers are also listed). The following are examples of port address text synonyms: bgp (179), dhcp (67), domain (53).
Directional	Refines entity type match conditions based on source and destination addressing. Directional match conditions are used only in conjunction with other match conditions.
destination	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
source	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
source and destination	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
source or destination	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	Matches packets according to packet size.
less *value*	Matches packets shorter than or equal to the specified value, in bytes.
greater *value*	Matches packets longer than or equal to the specified value, in bytes.
Protocol	Restricts matching condition to a particular protocol.
arp	Matches all ARP packets.
ether	Matches all Ethernet packets.
ether [*broadcast* | *multicast*]	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with source and destination
ether protocol [*address* | (*arp* | *ip* | *rarp*)]	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The ether protocol arguments arp, ip, and rarp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ether protocol match condition.
icmp	Matches all ICMP packets.
ip	Matches all IP packets.
ip [broadcast | *multicast*]	Matches broadcast or multicast IP packets.
ip protocol [*address* | (*icmp* | *igrp* | tcp | *udp*)]	Matches packets with the specified address or protocol type. The ip protocol arguments icmp, tcp, and udp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ip protocol match condition.
isis	Matches all IS-IS routing messages.
rarp	Matches all RARP packets.
tcp	Matches all TCP datagrams.
udp	Matches all UDP datagrams.

Table 54: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
||	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
( )	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors.

The arithmetic expression matching condition uses the following syntax:

monitor traffic matching "*arithmetic_expression relational_operator arithmetic_expression*"

The packet data accessor uses the following syntax:

protocol [byte-offset <*size*>]

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes

The following example captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

Table 55: Arithmetic and Relational Operators for the monitor traffic Command

Arithmetic or Relational Operator	Description
Arithmetic Operator	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
|	Bitwise inclusive OR.
Relational Operator (Highest to Lowest Precedence)	
< =	If the first expression is less than or equal to the second, the packet matches.
> =	If the first expression is greater than the second, the packet matches.
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.
=	If the compared expressions are equal, the packet matches.
! =	If the compared expressions are unequal, the packet matches.

Note

The syntax for the matching option is similar to the expression matching syntax in the UNIX tcpdump command. For extra tips and tricks for using monitor traffic, consult the tcpdump man page.

monitor interface

Syntax monitor interface <*interface-name* | traffic>

Description Display real-time statistics about a physical interface, updating them every second. The output of this command also shows the amount that each field has changed since you started the command or since you cleared the counters by using the C key. This command also checks for and displays common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.

To control the output of the command while it is running, use the following keys:

Table 56: Monitor Interface Output Control Keys

Action	Key
Display information about the next interface. The monitor interface command scrolls through the physical or logical interfaces in the same order that they are displayed by the show interfaces terse command.	N
Display information about a different interface. The command prompts you for the name of a specific interface.	I
Freeze the display, halting the display of updated statistics.	F
Thaw the display, resuming the display of updated statistics.	T
Clear (zero) the current delta counters since monitor interface was started. It does not clear the accumulative counter. To clear the accumulative counter, see clear interfaces statistics on page 193.	C
Stop the monitor interface command.	Q

Options *interface-name*—Name of a physical or logical interface.

traffic—Display traffic data for active interfaces.

Required Privilege Level trace

See Also show interfaces statistics on page 205 and monitor traffic on page 452

Output Fields router1—Host name of the router.

Seconds—How long the monitor interface command has been running or how long since you last zeroed the counters.

Time—Current time (UTC).

Interface—Short description of the interface, including its name, status, and encapsulation.

Link—State of the link. It can be Up, Down, or Test.

Current delta—Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last zeroed the counters.

Statistics—For an explanation of the interface statistics, see the description of the show interfaces statistics detail command for the appropriate interface type.

Sample Output: monitor interface (physical interface)

```
user@host> monitor interface so-0/0/0
router1                          Seconds: 19                  Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC48
Traffic statistics:                                          Current Delta
  Input packets:                     6045 (0 pps)                    [11]
  Input bytes:                    6290065 (0 bps)                [13882]
  Output packets:                   10376 (0 pps)                    [10]
  Output bytes:                  10365540 (0 bps)                 [9418]
Encapsulation statistics:
  Input keepalives:                  1901                             [2]
  Output keepalives:                 1901                             [2]
  NCP state: Opened
  LCP state: Opened
Error statistics:
  Input errors:                         0                             [0]
  Input drops:                          0                             [0]
  Input framing errors:                 0                             [0]
  Policed discards:                     0                             [0]
  L3 incompletes:                       0                             [0]
  L2 channel errors:                    0                             [0]
  L2 mismatch timeouts:                 0                             [0]
  Carrier transitions:                  1                             [0]
  Output errors:                        0                             [0]
  Output drops:                         0                             [0]
  Aged packets:                         0                             [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                             1                             [0]
  LOF count                             1                             [0]
  SEF count                             1                             [0]
  ES-S                                  0                             [0]
  SES-S                                 0                             [0]
SONET statistics:
  BIP-B1                           458871                             [0]
  BIP-B2                           460072                             [0]
  REI-L                            465610                             [0]
  BIP-B3                           458978                             [0]
  REI-P                            458773                             [0]
Received SONET overhead:
  F1      : 0x00  J0        : 0x00  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0x00
  C2(cmp) : 0x00  F2        : 0x00  Z3        : 0x00
  Z4      : 0x00  S1(cmp)   : 0x00
Transmitted SONET overhead:
  F1      : 0x00  J0        : 0x01  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0xcf
  F2      : 0x00  Z3        : 0x00  Z4        : 0x00

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

Sample Output: monitor interface (logical interface)

```
user@host> monitor interface so-1/0/0.0
host name                          Seconds: 16                  Time: 15:33:39
                                                                Delay: 0/0/1
Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:                                            Current delta
  Input bytes:                         0                              [0]
  Output bytes:                        0                              [0]
  Input packets:                       0                              [0]
  Output packets:                      0                              [0]
Remote statistics:
  Input bytes:                         0 (0 bps)                      [0]
  Output bytes:                        0 (0 bps)                      [0]
  Input packets:                       0 (0 pps)                      [0]
  Output packets:                      0 (0 pps)                      [0]
Traffic statistics:
  Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

Sample Output: monitor interface traffic

```
user@host> monitor interface traffic
host name                          Seconds: 15                  Time: 12:31:09

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-1/0/0    Down              0          (0)                  0          (0)
 so-1/1/0    Down              0          (0)                  0          (0)
 so-1/1/1    Down              0          (0)                  0          (0)
 so-1/1/2    Down              0          (0)                  0          (0)
 so-1/1/3    Down              0          (0)                  0          (0)
 t3-1/2/0    Down              0          (0)                  0          (0)
 t3-1/2/1    Down              0          (0)                  0          (0)
 t3-1/2/2    Down              0          (0)                  0          (0)
 t3-1/2/3    Down              0          (0)                  0          (0)
 so-2/0/0      Up         211035          (1)              36778          (0)
 so-2/0/1      Up         192753          (1)              36782          (0)
 so-2/0/2      Up         211020          (1)              36779          (0)
 so-2/0/3      Up         211029          (1)              36776          (0)
 so-2/1/0      Up         189378          (1)              36349          (0)
 so-2/1/1    Down              0          (0)              18747          (0)
 so-2/1/2    Down              0          (0)              16078          (0)
 so-2/1/3      Up              0          (0)              80338          (0)
 at-2/3/0      Up              0          (0)                  0          (0)
 at-2/3/1    Down              0          (0)                  0          (0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

monitor traffic

Syntax monitor traffic <absolute-sequence><count *count*><interface *interface*> <layer2-headers><matching *matching*><no-domain-names><no-optimize><no-resolve> <no-timestamp><print-ascii><print-hex><size *size*><brief | detail | extensive>

Description Print packet headers transmitted through network interfaces; these headers match a Boolean expression. The monitor traffic command is similar to the UNIX tcpdump command.

> **Caution** Using the monitor traffic command can negatively impact router performance. Not using print filtering functions, such as the count option or a matching expression, could impact packet throughput on your router.

See Also Monitor Traffic Match Conditions on page 445

Options absolute-sequence—(Optional) Print absolute TCP sequence numbers.

brief—(Optional) Display minimal protocol-related information. The monitor traffic command uses the brief option by default.

count *count*—(Optional) Specify the number of packet headers to display (0 through 1000000). The monitor traffic command quits automatically after displaying the number of packets specified.

detail—(Optional) Print packet data in moderate detail. For example, the detail option displays the time-to-live and type-of-service information in a TCP packet. For detailed output to be printed for some protocols, you must use the size option to increase the number of bytes of output that can be displayed for each packet matched by monitor traffic.

extensive—(Optional) Display the most extensive level of protocol-related output. For extensive output to be printed for certain protocols, you must use the size option to increase the output of each packet matched by monitor traffic. The extensive option displays link-level headers on each line.

interface *interface*—(Optional) Specify the interface on which monitor traffic prints packet data. If no interface is specified, monitor traffic prints packet data arriving on the lowest numbered interface.

layer2-headers—(Optional) Display the link-level header on each line.

matching *matching*—(Optional) Print packet headers that match a regular expression. Use matching expressions to define the level of detail with which monitor traffic filters and prints packet data.

For information about match conditions, see “Monitor Traffic Match Conditions” on page 445.

no-domain-names—(Optional) Suppress printing the domain portion of host names. With the no-domain-names option enabled, monitor traffic would print only team for the host name team.company.net.

no-resolve—(Optional) Suppress symbolic addressing.

no-timestamp—(Optional) Suppress timestamping on printed packets.

print-ascii—(Optional) Print each packet in ASCII format.

print-hex—(Optional) Print each packet, except for the link-level header, in hexadecimal format.

size *size*—(Optional) Receive the specified number of bytes for each packet. The default size is 68 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. The monitor traffic command truncates printed packets if the matched data exceeds the configured *size*.

Required Privilege Level

trace

Sample Output: monitor traffic

```
user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650  In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813  In host30.lab.home.net.syslog > host40.home0
18:17:28.817846  In host30.lab.home.net.syslog > host40.home0
18:17:28.819174  In host30.lab.home.net.syslog > host40.home0
...
```

Sample Output: monitor traffic (relative-sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp
count 20 matching "tcp"
listening on fxp0
 In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0) ack
1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39) ack 0 win 17680 (DF) [tos
0x10]  (ttl 64, id 5394)
 In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57) ack
1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

Sample Output: monitor traffic (absolute-sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp
count 20 matching "tcp" absolute-sequence
listening on fxp0
 In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0) ack
1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0]  (ttl )
 In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53) ack
1845421797 win 16384 <nop,nop,timestamp 4935628 965951>: BGP [|BGP UPDAT)
 In 192.168.4.227.1024 > 207.17.136.193.179: P 1845421797:1845421852(55) ack
4042780912 win 16384 <nop,nop,timestamp 965951 4935628>: BGP [|BGP UPDAT)
...
```

Sample Output: monitor traffic (count)

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125  In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185 Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

Sample Output: monitor traffic detail

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864  In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926 Out my-server.work.net.telnet.telnet > my-server.home.net.1295:
P 1:38(37) ack 0 win 17680 (DF) [tos 0x10]  (ttl 6)
```

Sample Output: monitor traffic extensive

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp0
13:18:17.406933  In 192.168.4.206.2723610880 > 172.17.28.8.2049: 40 null (ttl
64, id 38367)
13:18:17.407577  In 172.17.28.8.2049 > 192.168.4.206.2723610880: reply ok 28
null (ttl 61, id 35495)
13:18:17.541140  In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
                         0000 0100 0000 0000 0000 0000 0000 0000
                         0000 0000 0000 0000 0000 0000 0000 0000
                         0000 0000 0000 0000 0000 0000 0000
13:18:17.591513  In 172.24.248.156.4139 > 192.168.4.210.23: .
3556964918:3556964918(0) ack 295526518 win 17601 (DF) (ttl 121, id 14)
13:18:17.591568 Out 192.168.4.210.23 > 172.24.248.156.4139: P 1:40(39) ack 0 win
17680 (DF) [tos 0x10]  (ttl 64, id 52376)
```

show aps

Syntax show aps <interface *interface-name* | group *group*> <brief | detail | extensive | summary>

Description Display information about Automatic Protection Switching (APS).

Options brief—(Optional) Display brief APS information.

detail—(Optional) Display detailed APS information.

extensive—(Optional) Display very detailed APS information.

summary—(Optional) Display summary APS information.

interface *interface-name*—(Optional) Display APS information for the named interface.

group *group*—(Optional) Display APS information for the named group.

Default: brief

Required Privilege Level view

Sample Output
Sample Output: show aps (standard) on page 457
Sample Output: show aps brief on page 457
Sample Output: show aps detail on page 457
Sample Output: show aps extensive on page 457
Sample Output: show aps summary on page 457
Sample Output: show aps group on page 458
Sample Output: show aps interface on page 458

Options at a Glance Table 57 summarizes which information is included in the output of each show aps command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 57: Show APS Output Field Summary

Options	Field Description
Detail Extensive	adj—State of the neighbor adjacency. It can be Up, Down, Init, or Unknown.
Detail Extensive	Channel state—Displays which circuit has been selected. It can be Working or Protect.
All	Circuit—Circuit type. It can be Working or Protect.
Detail Extensive	dead—Number of seconds before neighbor is declared dead.
All	Group—Group name.
Extensive	Hello due—Time next hello packet will be sent.
All	Interface—Name of interface.
All	Intf State—State of the interface and circuit. The interface can be Disabled or Enabled. The circuit can be Up, Degraded, Down, Unknown, or Non-existent.
Extensive	nbr K1— Displays the value of the SONET/SDH K1 byte requested to be transmitted by the neighbor.
Extensive	nbr paired req—Nonzero if the neighbor is requesting a particular K1 value due to a change in the paired circuit.
Extensive	neighbor revert time—Configured time, on the neighbor interface, to wait after the working circuit has again become functional before making the working circuit active again.

Options	Field Description
Detail Extensive	Neighbor—Address and state of neighbor interface. If the working and protect interfaces are on the same router, the neighbor address is displayed as 0.0.0.0.
Detail Extensive	Protect circuit is on—When both the working circuit and protect circuit are on the same router, displays interface name of APS protect circuit.
Extensive	rcv K1—Displays the value of the SONET/SDH K1 byte received on this interface. (Valid only on the protect circuit.)
Extensive	Req K1—Displays the value of the SONET/SDH K1 byte requested to be transmitted by this circuit.
Extensive	Revert time—Configured time to wait after the working circuit has again become functional before making the working circuit active again.
Detail Extensive	Working circuit is on—When both the working circuit and protect circuit are on the same router, displays interface name of APS working circuit.
Extensive	xmit K1—Displays the value of the SONET/SDH K1 byte being transmitted on this interface. (Valid only on the protect circuit.)

Output Fields

Interface—Name of interface.

Group—Group name.

Circuit—Circuit type. It can be Working or Protect.

Intf State—State of the interface and circuit. The interface can be Disabled or Enabled. The circuit can be Up, Degraded, Down, Unknown, or Non-existent.

Neighbor—(Detail and extensive output only) Address and state of neighbor interface. If the working and protect interfaces are on the same router, the neighbor address is displayed as 0.0.0.0.

adj—(Detail and extensive output only) State of the neighbor adjacency. It can be Up, Down, Init, or Unknown.

dead—(Detail and extensive output only) Number of seconds before neighbor is declared dead.

Channel state—(Detail and extensive output only) Display which circuit has been selected. It can be Working or Protect.

Protect circuit is on—(Detail and extensive output only) When both the working circuit and protect circuit are on the same router, displays interface name of APS protect circuit.

Working circuit is on—(Detail and extensive output only) When both the working circuit and protect circuit are on the same router, displays interface name of APS working circuit.

Req K1—(Extensive output only) Display the value of the SONET/SDH K1 byte requested to be transmitted by this circuit.

rcv K1—(Extensive output only) Display the value of the SONET/SDH K1 byte received on this interface. (Valid only on the protect circuit.)

xmit K1—(Extensive output only) Display the value of the SONET/SDH K1 byte being transmitted on this interface. (Valid only on the protect circuit.)

nbr K1—(Extensive output only) Display the value of the SONET/SDH K1 byte requested to be transmitted by the neighbor.

nbr paired req—(Extensive output only) Nonzero if the neighbor is requesting a particular K1 value due to a change in the paired circuit.

Revert time—(Extensive output only) Configured time to wait after the working circuit has again become functional before making the working circuit active again.

neighbor revert time—(Extensive output only) Configured time, on the neighbor interface, to wait after the working circuit has again become functional before making the working circuit active again.

Hello due—(Extensive output only) Time next hello packet will be sent.

Sample Output: show aps (standard)

```
user@host> show aps
Interface  Group                          Circuit  Intf State
t3-1/2/0:0 king                           Working  enabled, up
t3-1/3/0:0 king                           Protect  disabled, up
```

Sample Output: show aps brief

```
user@host> show aps brief
Interface  Group                          Circuit  Intf State
t3-1/2/0:0 king                           Working  enabled, up
t3-1/3/0:0 king                           Protect  disabled, up
```

Sample Output: show aps detail

```
user@host> show aps detail
Interface  Group                          Circuit  Intf State
t3-1/2/0:0 king                           Working  enabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface disabled, dead 2.670
  Channel state Working
  Protect circuit is on interface t3-1/3/0:0
t3-1/3/0:0 king                           Protect  disabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface enabled, dead 2.750
  Channel state Working
  Working circuit is on interface t3-1/2/0:0
```

Sample Output: show aps extensive

```
user@host> show aps extensive
Interface  Group                          Circuit  Intf State
t3-1/2/0:0 king                           Working  enabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface disabled, dead 2.100
  Channel state Working
  Protect circuit is on interface t3-1/3/0:0
  Req K1 0x00, rcv K1 0x00, xmit K1 0x00, nbr K1 0x11, nbr paired req 0
  Revert time 0, neighbor revert time 0
  Hello due in 0.258
t3-1/3/0:0 king                           Protect  disabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface enabled, dead 2.400
  Channel state Working
  Working circuit is on interface t3-1/2/0:0
  Req K1 0x11, rcv K1 0x00, xmit K1 0x00, nbr K1 0x00, nbr paired req 0
  Revert time 0, neighbor revert time 0
  Hello due in 0.076
```

Sample Output: show aps summary

```
user@host show aps summary
Interface  Group                          Circuit  Intf State
t3-1/2/0:0 king                           Working  enabled, up
t3-1/3/0:0 king                           Protect  disabled, u
```

Sample Output: show aps group

```
user@host> show aps group king
Interface  Group                          Circuit  Intf State
t3-1/3/0:0 king                           Protect  disabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface enabled, dead 2.900
  Channel state Working
  Working circuit is on interface t3-1/2/0:0
  Req K1 0x11, rcv K1 0x00, xmit K1 0x00, nbr K1 0x00, nbr paired req 0
  Revert time 0, neighbor revert time 0
  Hello due in 0.633
t3-1/2/0:0 king                           Working  enabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface disabled, dead 2.860
  Channel state Working
  Protect circuit is on interface t3-1/3/0:0
  Req K1 0x00, rcv K1 0x00, xmit K1 0x00, nbr K1 0x11, nbr paired req 0
  Revert time 0, neighbor revert time 0
  Hello due in 0.836
```

Sample Output: show aps interface

```
user@host> show aps interface t3-1/3/0:0
Interface  Group                          Circuit  Intf State
t3-1/3/0:0 king                           Protect  disabled, up
  Neighbor 0.0.0.0, adj up, neighbor interface enabled, dead 2.230
  Channel state Working
  Working circuit is on interface t3-1/2/0:0
  Req K1 0x11, rcv K1 0x00, xmit K1 0x00, nbr K1 0x00, nbr paired req 0
  Revert time 0, neighbor revert time 0
  Hello due in 0.491
```

Chapter 28
Monitoring and Troubleshooting using BERT Tests

Table 58 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot BERT tests on router interfaces. In the table, the commands are grouped by functionality. In the remainder of this chapter, in order of appearance.

Table 58: Commands for Monitoring Interfaces Using BERT Tests

Task or Information to Monitor	CLI Command
Start bit error rate testing.	test interface bert-start on page 459
Stop bit error rate testing.	test interface bert-stop on page 460
Transmit over FDL to initiate or terminate a far-end line loopback.	test interface feac-loop-initiate on page 463
Transmit over FDL to initiate or terminate a far-end payload loopback.	test interface fdl-payload-loop on page 462
Transmit the line loopback activate code word sequence on the interface's FEAC channel.	test interface feac-loop-initiate on page 463
Transmit the line loopback deactivate code word sequence on the interface's FEAC channel.	test interface feac-loop-terminate on page 463

test interface bert-start

Syntax test interface t3-*fpc/pic/port <:channel>* bert-start

Description Start bit error rate testing. To start a BERT test, the interface must be disabled. You can run a BERT test on only one interface per PIC at any one time.

To configure the BERT parameters—the algorithm, bit error rate, and test duration—include the bert-algorithm, bert-error-rate, and bert-period statements, respectively, at the [edit interfaces *interface-name* t3-options] hierarchy level.

To display the results of the BERT test, use the show interfaces extensive command.

Options t3-*fpc/pic/port <:channel>*—Name of interface. Channel number indicates Channelized OC-12 interface.

Sample Output

```
user@host> test interface t3-5/2/0 bert-start
user@host> show interfaces t3-5/2/0 extensive
...
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
    BERT time period: 10 seconds, Elapsed: 10 seconds (completed)
    Algorithm: 2^3 - 1, Pseudorandom (1), Error rate: 10e-0
    Bit count      :         5656, Overflows: 0
    Error bit count:          811, Overflows: 0
    LOS status: OK, LOS count: 1, LOS seconds: 9
...
```

Required Privilege Level view

See Also show interfaces on page 194, test interface bert-stop on page 460

test interface bert-stop

Syntax test interface t3-*fpc/pic/port <:channel>* bert-stop

Description Stop bit error rate testing.

To display the results of the BERT test, use the show interfaces extensive command.

Options t3-*fpc/pic/port <:channel>*—Name of a T3 interface. Channel number indicates Channelized OC-12 interface.

Required Privilege Level view

See Also show interfaces on page 194, test interface bert-start on page 459

test interface fdl-line-loop

Syntax test interface t1-*fpc/pic/port <:channel>* fdl-line-loop ansi| bellcore initiate|terminate

Description Send commands over the Facility Data Link (FDL) on a T1 interface to initiate or terminate a far-end line loopback using either an ANSI or Bellcore FDL command code. If the far end of the connection is in C-bit parity mode and it has been configured to accept line loopback requests from the near end, the far end executes the request.

> **Note** Juniper Networks T1 interfaces do not respond to this command. For more information, see the ANSI T1.107 specification.

Options t1-*fpc/pic/port <:channel>*—Name of a T1 interface. Channel number indicates Channelized DS3 to DS1 interface.

ansi—ANSI FDL command code.

bellcore—Bellcore FDL command code.

intitate—Initiate the far-end payload loopback.

terminate— Terminate the far-end payload loopback.

Required Privilege Level view

See Also test interface fdl-payload-loop on page 462

test interface fdl-payload-loop

Syntax test interface t1-*fpc/pic/port <:channel>* fdl-payload-loop ansi| bellcore initiate|terminate

Description Send commands over the Facility Data Link (FDL) on a T1 interface to initiate or terminate a far-end payload loopback using either an ANSI or Bellcore FDL command code. If the far end of the connection is in C-bit parity mode and has been configured to accept payload loopback requests from the near end, the far end executes the request.

> **Note** Juniper Networks T1 interfaces do not respond to this command. For more information, see the ANSI T1.107 specification.

Options t1-*fpc/pic/port <:channel>*—Name of a T1 interface. Channel number indicates Channelized DS3 to DS1 interface.

ansi—ANSI FDL (Facility Data Link) command code.

bellcore—Bellcore FDL command code.

intitate—Initiate the far-end payload loopback.

terminate—Terminate the far-end payload loopback.

Required Privilege Level view

See Also test interface feac-loop-initiate on page 463

test interface feac-loop-initiate

Syntax test interface t3-*fpc/pic/port <:channel>* feac-loop-initiate

Description Have the interface transmit the word sequence for the line loopback activate code on its FEAC channel. If the far end of the connection is in C-bit parity mode and has been configured to accept remote loopback requests from the near end, the far end places its interface into remote loopback. See the ANSI T1.107 specification for more details.

If the far end is a Juniper Networks router, it accepts the request and places itself into remote loopback if the following are true:

- The router is configured to respond to FEAC loop requests; that is, the feac-loop-respond statement is included at the [edit interfaces *interface-name* t3-options] hierarchy level.
- The router is running in C-bit parity mode (enabled by default).
- The router has no locally configured local or remote loopback; that is, no loopback statement is included at the [edit interfaces *interface-name* t3-options] hierarchy level.

To display the state of FEAC-initiated loopback and the number of times the interface has placed itself into remote loopback using a FEAC request, use the show interfaces extensive command.

Options t3-*fpc/pic/port <:channel>*—Name of a T3 interface. Channel number indicates Channelized OC-12 interface.

Required Privilege Level view

See Also show interfaces on page 194, test interface feac-loop-terminate on page 463; for information about the feac-loop-respond statement and the loopback statement, see the *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service.*

test interface feac-loop-terminate

Syntax test interface t3-*fpc/pic/port <:channel>* feac-loop-terminate

Description Have the interface transmit the line loopback deactivate code word sequence on its FEAC channel. If the far end of the connection is in C-bit parity mode and has been configured to accept remote loopback requests from the near end, the far end clears remote loopback on the interface. See the ANSI T1.107 specification for more details.

If the far end is a Juniper Networks router that has been configured to respond to FEAC loop requests (that is, the feac-loop respond statement is included at the [edit interfaces *interface-name* t3-options] hierarchy level), is in C-bit parity mode (enabled by default), and has no locally configured local or remote loopback (that is, no loopback statement is included at the [edit interfaces *interface-name* t3-options] hierarchy level), the far-end interface accepts the request and clears remote loopback.

Options t3-*fpc/pic/port <:channel>*—Name of a T3 interface. Channel number indicates Channelized OC-12 interface.

Required Privilege Level view

See Also show interfaces on page 194, test interface feac-loop-initiate on page 463

Chapter 29

VRRP Router Interfaces Monitoring and Troubleshooting

Table 59 summarizes the command-line interface (CLI) commands you can use to monitor land troubleshoot VRRP router interfaces. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 59: Table 59 Commands for Monitoring VRRP Interfaces

Task or Information to Monitor	CLI Command
VRRP groups.	show vrrp on page 466
Clear (zero) VRRP groups.	clear vrrp on page 465

clear vrrp

Syntax clear vrrp statistics (*interface-name* | all)

Description Zero interface statistics.

Options all—Zero statistics on all interfaces.

interface-name—Name of the interface.

Required Privilege Level clear

See Also show vrrp on page 466

show vrrp

Syntax show vrrp <brief | detail | extensive | summary> (track | interface *interface-name*)

Description Display information and status about VRRP groups.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

extensive—(Optional) Display very detailed information.

interface *interface-name*—(Optional) Name of an interface.

summary—(Optional) Display very brief information.

track—Display information and status about VRRP groups.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show vrrp brief on page 470
Sample Output: show vrrp detail on page 471
Sample Output: show vrrp extensive on page 471
Sample Output: show vrrp summary on page 470

The sample outputs shown in the following sections are for the following configurations:

Master router:

```
[edit]
user@master-router# show interfaces ge-3/1/0
unit 0
   family-inet {
      address 192.168.29.254/24 {
         vrrp-group 10 {
            virtual-address 192.168.29.55;
            priority 200;
            advertise-interval 3;
            authentication-type simple;
            authentication-key “booJUM”;
         }
      }
   }
}
```

Backup router:

```
[edit]
user@backup-router# show interfaces ge-5/2/0
unit 0
   family-inet {
      address 192.168.29.10/24 {
         vrrp-group 10 {
            virtual-address 192.168.29.55;
            priority 190;
            advertise-interval 3;
            authentication-type simple;
            authentication-key "booJUM";
         }
      }
   }
}
```

Options at a Glance Table 60 summarizes which information is included in the output of each show vrrp command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 60: Show VRRP Output Field Summary

Options	Field Description
Extensive	Active—Total number of VRRP groups that are active (that is, whose interface state is either up or down).
Detail Extensive	Address—Address of the physical interface.
Detail Extensive	Advertisement interval—Configured VRRP advertisement interval.
Detail Extensive	Advertisement timer—How long until the advertisement timer expires, in seconds.
Detail Extensive	Authentication type—Configured VRRP authentication type. It can be none, simple, or md5.
Detail Extensive	Dead timer—How long until the Master Is Dead timer expires, in seconds.
Extensive	Group state transition statistics—State transition statistics for the VRRP group.
Extensive	Group VRRP PDU error statistics—Errored statistics for the VRRP group.
Extensive	Group VRRP PDU statistics—Number of VRRP advertisements sent and received by the group.
Extensive	Groups—Total number of VRRP groups configured on the interface.
All	Group—VRRP group number.
All	Int state—State of the physical interface.
Extensive	Interface index—Physical interface's index number, which reflects its initialization sequence.
Extensive	Interface VRRP PDU error statistics—Errored statistics for the logical interface.
Extensive	Interface VRRP PDU statistics—Nonerrored statistics for the logical interface.
All	Interface—Name of the local interface.
Detail Extensive	Master priority—Priority value of the router acting as the master.
Detail Extensive	Master router—IP address of the interface that is acting as the master.
Detail Extensive	Physical interface—Name of the interface.
Detail Extensive	Preempt—Whether preemption is allowed on the interface.
Detail Extensive	Priority—Configured VRRP priority for the interface.
Detail Extensive	State—VRRP state information.

Options	Field Description
Brief	Timer—VRRP timer information: ■ A—How long until the advertisement timer expires, in seconds. ■ D—How long until the Master Is Dead timer expires, in seconds.
Brief Summary	Type and Address—Identifier for the address and the address itself.
All	Unit—Logical unit number.
Detail Extensive	VIP count—Number of virtual IP addresses that have been configured on the interface.
Detail Extensive	VIP—List of virtual IP addresses configured on the interface.
Detail Extensive	Virtual MAC—MAC address associated with the virtual IP address.
Brief Summary	VR state—VRRP state information.

Output Fields

Interface—Name of the logical interface.

Interface index—(Extensive output only) Physical interface's index number, which reflects its initialization sequence.

Groups—(Extensive output only) Total number of VRRP groups configured on the interface.

Active—(Extensive output only) Total number of VRRP groups that are active (that is, whose interface state is either up or down).

Interface VRRP PDU statistics—(Extensive output only) Nonerrored statistics for the logical interface:

- Advertisement sent—Number of VRRP advertisement PDUs that the interface has transmitted.
- Advertisement received—Number of VRRP advertisement PDUs received from an interface.
- Packets received—Number of VRRP packets received for VRRP groups on the interface.
- No group match received—Number of VRRP packets received for VRRP groups that do not exist on the interface.

Interface VRRP PDU error statistics—(Extensive output only) Errored statistics for the logical interface:

- Invalid IPAH next type received—Number of packets received that use IPAH and that do not encapsulate VRRP packets.
- Invalid VRRP TTL value received—Number of packets received whose IP TTL value is not 255.
- Invalid VRRP version received—Number of packets received whose VRRP version is not 2.
- Invalid VRRP PDU type received—Number of packets received whose VRRP PDU type is not 1.
- Invalid VRRP authentication type received—Number of packets received whose VRRP authentication is not none, simple, or md5.

- Invalid VRRP IP count received—Number of packets received whose VRRP IP count exceeds 8.
- Invalid VRRP checksum received—Number of packets received whose VRRP checksum does not match the calculated one.

Physical interface—Name of the physical interface.

Unit—Logical unit number.

Address—Address of the physical interface.

Type and Address—Identifier for the address and the address itself:

- lcl—Configured local interface address.
- mas—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router.
- vip—Configured virtual IP addresses.

Interface state—State of the physical interface. It can be one of the following:

- down—The device is present and the link is down.
- not present—The interface is configured, but no physical device is present.
- unknown—The VRRP process has not had time to query the kernel about the state of the interface.
- up—The device is present and the link is up.

Group—VRRP group number.

State—VRRP state information. It can be one of the following:

- backup—The interface is acting as the backup router interface.
- bringup—VRRP is just starting, and the physical device is not yet present.
- idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is up.
- initializing—VRRP is initializing.
- master—The interface is acting as the master router interface.
- transition—The interface is changing between being the backup and the master router.

Priority—Configured VRRP priority for the interface.

Advertisement interval—Configured VRRP advertisement interval.

Authentication type—Configured VRRP authentication type. It can be none, simple, or md5.

Preempt—Whether preemption is allowed on the interface. It can be yes or no.

VIP count—Number of virtual IP addresses that have been configured on the interface.

VIP—List of virtual IP addresses configured on the interface.

Advertisement timer—How long until the advertisement timer expires, in seconds.

Master router—IP address of the interface that is acting as the master.

Virtual MAC—MAC address associated with the virtual IP address.

Dead timer—How long until the Master Is Dead Timer expires, in seconds.

Master priority—Priority value of the router acting as the master.

Group VRRP PDU statistics—(Extensive output only) Number of VRRP advertisements sent and received by the group.

Group VRRP PDU error statistics—(Extensive output only) Errored statistics for the VRRP group. For an explanation of the individual fields, see the explanation of the Interface VRRP PDU statistics field above.

Group state transition statistics—(Extensive output only) State transition statistics for the VRRP group.

VR state—VRRP information. It can be one of the following:

- backup—The interface is acting as the backup router interface.
- bringup—VRRP is just starting, and the physical device is not yet present.
- idle—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is up.
- initializing—VRRP is initializing.
- master—The interface is acting as the master router interface.
- transition—The interface is changing between being the backup and the master router.

Timer—(Brief output only) VRRP timer information:

- A—How long until the advertisement timer expires, in seconds.
- D—How long until the Master Is Dead timer expires, in seconds.

Sample Output: show vrrp brief

```
user@backup-router> show vrrp brief
Interface   Unit  Group  Type  Address          Int state    VR state   Timer
ge-5/2/0    0     10     lcl   192.168.29.10    up           backup     D 7.133
                         vip   192.168.29.55
                         mas   192.168.29.254
```

Sample Output: show vrrp summary

```
user@backup-router> show vrrp summary
Interface   Unit  Group  Type  Address          Int state    VR state
ge-5/2/0    0     10     lcl   192.168.29.10    up           backup
                         vip   192.168.29.55
                         mas   192.168.29.254
```

Sample Output: show vrrp interface

```
user@master-router> show vrrp interface ge-3/1/0
Interface: ge-3/1/0, Ifindex: 1, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          377
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP ttl value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP pdu type received          :            0
    Invalid VRRP authentication type received:            0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical Interface: ge-3/1/0, Unit: 0, Address: 192.168.29.254/24
  Interface state: up, Group: 10, State: master
  Priority: 200, Advertisement interval: 3, Authentication type: simple
  Preempt: yes, VIP count: 1, VIP: 192.168.29.55
  Advertisement timer : 1.597s, Master router: 192.168.29.254
  Virtual MAC: 00:00:5e:00:01:0a
  Group vrrp pdu statistics
    Advertisement sent              :          377
    Advertisement received          :            0
  Group vrrp pdu error statistics
    Bad auth type received          :            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
   Idle to master transitions       :            0
   Idle to backup transitions       :            1
   Backup to master transitions     :            1
   Master to backup transitions     :            0
```

Sample Output: show vrrp detail

```
user@master-router> show vrrp detail
Physical interface: ge-3/1/0, Unit: 0, Address: 192.168.29.254/24
  Interface state: up, Group: 10, State: master
  Priority: 200, Advertisement interval: 3, Authentication type: simple
  Preempt: yes, VIP count: 1, VIP: 192.168.29.55
  Advertisement timer: 2.407s, Master router: 192.168.29.254
  Virtual MAC: 00:00:5e:00:01:0a
```

Sample Output: show vrrp extensive

```
user@master-router> show vrrp extensive
Interface: ge-3/1/0, Interface index: 1, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          377
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP ttl value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP pdu type received          :            0
    Invalid VRRP authentication type received:            0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0
```

```
Physical Interface: ge-3/1/0, Unit: 0, Address: 192.168.29.254/24
  Interface state: up, Group: 10, State: master
  Priority: 200, Advertisement interval: 3, Authentication type: simple
  Preempt: yes, VIP count: 1, VIP: 192.168.29.55
  Advertisement timer : 1.597s, Master router: 192.168.29.254
  Virtual MAC: 00:00:5e:00:01:0a
  Group VRRP PDU statistics
    Advertisement sent              :          377
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
   Idle to master transitions       :            0
   Idle to backup transitions       :            1
   Backup to master transitions     :            1
   Master to backup transitions     :            0
t3-5/2/0 extensive
...
  DSU configuration:
    Compatibility mode: None, Scrambling: Disabled, Subrate: Disabled
    FEAC loopback: Inactive, Response: Disabled, Count: 0
    BERT time period: 10 seconds, Elapsed: 10 seconds (completed)
    Algorithm: 2^3 - 1, Pseudorandom (1), Error rate: 10e-0
    Bit count      :         5656, Overflows: 0
    Error bit count:          811, Overflows: 0
    LOS status: OK, LOS count: 1, LOS seconds: 9
```

Part 4
Policy Framework

- Firewall Filters Monitoring and Troubleshooting on page 475
- Forwarding Options Monitoring and Troubleshooting on page 479
- Routing Policy Monitoring and Troubleshooting on page 483

Chapter 30
Firewall Filters Monitoring and Troubleshooting

Table 61 summarizes the command-line interface (CLI) commands you can use to monitor firewall filters. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 61: Commands for Monitoring and Managing Firewall Filters

Task or Information to Monitor	Command
Counters for firewall filters	show firewall filter on page 476
Clear firewall filter counters	clear firewall on page 475
Firewall filter log information	show firewall log on page 477

clear firewall

Syntax clear firewall (all | counter *counter-name* | filter *filter-name*)

Description Clear statistics about configured firewall filters.

Options all—Clear the packet and byte counts for all filters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

Required Privilege Level clear

show firewall filter

Syntax show firewall filter *filter-name*

Description Display statistics about configured firewall filters.

Options *filter-name*—Display information about a specific filter.

Required Privilege Level view

Output Fields Filter—Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level. When an interface-specific filter is used, the name of the filter is followed by the full interface name and by either -i indicating an input filter, or -o indicating an output filter.

Counter—Name of a filter counter that has been configured with the counter firewall filter action.

Packet count—Number of packets that matched the filter.

Byte count—Number of bytes that matched the filter.

Sample Output

```
user@host> show firewall filter hello1
Filter/Counter                      Packet count            Byte count
ip-option-filter
  rr-count                                    27                  3348
  ts-count                                     0                     0
  ssrr-count                                   0                     0
  lsrr-count                                   0                     0
  ra-count                                 11654               2493120
```

The following example shows output from an interface specific firewall filter. bob-ge-1/0/0.0-i indicates an input filter. bob-ge-1/0/0.0-o indicates an output filter.

```
user@host> show firewall filter
Filter/Counter                      Packet count            Byte count
ip-option-filter
  liz                                        372                 54831
f3-ge-1/0/0.0-i
   bob-ge-1/0/0.0-i                           49                  4100
f3-ge-1/0/0.0-o
   bob-ge-1/0/0.0-o                           56                  3844
```

show firewall log

Syntax show firewall log <interface *interface-name*> <detail>

Description Display log information about firewall filters.

Options detail—(Optional) Display detailed information.

interface *interface-name*—(Optional) Display log information about a specific interface.

Required Privilege Level view

Output Fields Time—Time that the event occurred.

Filter—Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level. A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine. A space (no hyphen) indicates the packet was handled by the Routing Engine. pfe indicates packets logged from the PFE hardware filters.

Act—Filter action:

- A—Accept
- D—Discard
- R—Reject

Interface—Ingress interface for the packet.

Pro—Packet's protocol name or number. It can be one of the following: egp, gre, icmp, ipip, ospf, pim, rsvp, tcp, or udp.

Source address—Packet's source address.

Destination address—Packet's destination address and port.

Sample Output

```
user@host> show firewall log
Time     Filter     A Interface        Pro Source address  Destination address
16:08:04 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:24373
16:08:03 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:29531
16:08:02 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:27265
16:08:01 pfe        A so-1/1/0.0       OSP 123.168.10.65   212.0.0.5:48
16:08:01 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:43943
16:08:00 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:58572
16:07:59 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:56307
16:07:58 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:60185
16:07:57 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:1600
16:07:56 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:6502
16:07:55 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:17548
16:07:54 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:5298
16:07:53 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:24536
16:07:52 sample-test A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:24373
16:07:52 sample-test A local            ICM 123.168.10.66   123.168.10.65:22325
16:07:52 pfe        A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:27900
16:07:51 pfe        A so-1/1/0.0       OSP 123.168.10.65   212.0.0.5:48
16:07:51 sample-test A so-1/1/0.0       ICM 123.168.10.65   123.168.10.66:29531
16:07:51 sample-test A local            ICM 123.168.10.66 123.168.10.65:27483
...
```

Chapter 31 Forwarding Options Monitoring and Troubleshooting

Table 62 summarizes the command-line interface (CLI) commands you can use to monitor forwarding options. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 62: Commands for Forwarding Options Monitoring and Troubleshooting

Task Category	Task or Information to Monitor	Command
Forwarding Options Information	Show statistics collected by the UDP forwarding process.	show helper statistics on page 480
	Clear statistic counters in the UDP forwarding process.	clear helper statistics on page 479
	Restart the UDP forwarding process.	restart helper on page 481

clear helper statistics

Syntax clear helper statistics

Description Clear statistic counters in the User Datagram Protocol (UDP) forwarding process.

Required Privilege Level view

Sample Output

```
user@host> show helper statistics
domain:
  Received packets: 63
  Forwarded packets: 61
  Dropped packets: 2
    Due to no interface in fud database: 0
    Due to an error during packet read: 1
    Due to an error during packet send: 1
tftp:
  Received packets: 5
  Forwarded packets: 5
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0

user@host>clear helper statistics
```

```
user@host>show helper statistics
domain:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
tftp:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
```

show helper statistics

Syntax show helper statistics

Description Show statistics collected by the User Datagram Protocol (UDP) forwarding process.

Required Privilege Level view

Output Fields Domain—Statistics for DNS forwarding.

Received packets—Packets received for this service.

Forwarded packets—Packets forwarded for this service.

Dropped packets—Total number of packets dropped for this service.

Due to no interface in fud database—Number of packets dropped because the packet came in on an interface that the UDP forwarding process was not expecting to be active.

Due to an error during packet read—Number of packets dropped because there was an error when the packet was read from the wire.

Due to an error during packet send—Number of packets dropped because there was an error when the packet was sent to the wire.

tftp—Statistics for TFTP forwarding.

Sample Output

```
user@host> show helper statistics
domain:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
tftp:
  Received packets: 0
  Forwarded packets: 0
  Dropped packets: 0
    Due to no interface in fud database: 0
    Due to an error during packet read: 0
    Due to an error during packet send: 0
```

restart helper

Syntax restart helper

Description Restart the User Datagram Protocol (UDP) forwarding process.

Required Privilege Level view

Sample Output

```
user@host> restart helper
Forwarding UDP daemon started, pid 9261
```

Chapter 32
Routing Policy Monitoring and Troubleshooting

Table 63 summarizes the command-line interface (CLI) commands you can use to monitor the routing policy filters. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 63: Commands for Monitoring and Managing Routing Policy

Task or Information to Monitor	Command
Test import and export policies.	test policy on page 484
Configured routing policies.	show policy on page 483
Configured BGP damping parameters that can be used in policy filters.	show policy damping on page 484
Display routes that have been damped.	show route inactive on page 549

show policy

Syntax show policy <*policy-name*>

Description Display information about the configured routing policies.

Options none—List the names of all configured routing policies.

policy-name—(Optional) Show the contents of the specified policy.

Required Privilege Level view

Sample Output

```
user@host> show policy
Configured policies:
test-statics
> show policy test-statics
Policy test-statics:
    from
         3.0.0.0/8  accept
         3.1.0.0/16  accept
    then reject
```

show policy damping

Syntax show policy damping

Description Display information about the BGP damping parameters.

Required Privilege Level view

Output Fields Halflife—Decay half-life. minutes is the time duration during which the accumulated figure-of-merit value is reduced by half if the route remains stable. The default is 15 minutes.

Maximum suppress time—Maximum hold-down time. minutes is the maximum time that a route can be suppressed no matter how unstable it has been prior to this period of stability. The default is 60 minutes.

Sample Output

```
user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
```

test policy

Syntax test policy *policy-name prefix*

Description Test a policy configuration to determine which prefixes match routes in the routing table. show policy damping knows which interface the routes are coming from.

All prefixes in the inet.0 routing table (the default unicast routing table) that match the specified prefix (that is, prefixes that are the same or are longer matches) are processed by the from clause in the specified policy. All prefixes accepted by the policy are displayed. show policy damping does have the nexthop interface available.

Options *policy-name*—Name of a policy. To test all policies, specify *policy-name* as a hyphen (–).

prefix—Destination prefix to match.

Required Privilege Level view

Sample Output

```
user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

3.0.0.0/8          *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
                   AS Path: 50888 I
                    > to 10.11.4.32 via en0.2, label-switched-path l2
3.3.3.1/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                    > to 10.0.4.7 via fxp0.0
3.3.3.2/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                    > to 10.0.4.7 via fxp0.0
3.3.3.3/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                    > to 10.0.4.7 via fxp0.0
3.3.3.4/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                    > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

Part 5
Class of Service

■ Class of Service Monitoring and Troubleshooting on page 487

Chapter 33

Class of Service Monitoring and Troubleshooting

Table 64 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot Class of Service (CoS). In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 64: Commands for Monitoring and Troubleshooting Class of Service

	Task or Information to Monitor	CLI Command
Commands to query CoS statistics	Interface queue information	show interfaces queue on page 506
	Queue statistics summary per interface	show interfaces extensive (CoS related attributes only) on page 504
Commands to query CoS configuration	Extensive CoS configuration information	show class-of-service on page 488
	Mapping of code point value to forwarding class and loss priority	show class-of-service classifier on page 489
	Code point aliases	show class-of-service code-point-aliases on page 490
	Random Early Detection (RED) profiles	show class-of-service drop-profile on page 491
	Forwarding class to queue mapping	show class-of-service forwarding-class on page 492
	Per interface CoS configuration	show class-of-service interface on page 500
	Rewrite rules	show class-of-service rewrite-rule on page 501
	Scheduler to forwarding class mapping	show class-of-service scheduler-map on page 502

	Task or Information to Monitor	CLI Command
Commands to query CoS configuration in the forwarding table	CoS configuration as it exists in the forwarding table	show class-of-service forwarding-table on page 492
	Classifiers in the forwarding table	show class-of-service forwarding-table classifier on page 493
	Per logical interface classifier information in the forwarding table	show class-of-service forwarding-table classifier mapping on page 494
	Per physical interface scheduler information in the forwarding table	show class-of-service forwarding-table scheduler-map on page 498
	Random Early Detection (RED) profiles in the forwarding table	show class-of-service forwarding-table drop-profile on page 495
	Rewrite rules in the forwarding table	show class-of-service forwarding-table rewrite-rule on page 496
	Per logical interface rewrite rule information in the forwarding table	show class-of-service forwarding-table rewrite-rule mapping on page 497

show class-of-service

Syntax show class-of-service

Description Display entire CoS configuration, including system-chosen defaults.

Executing this command is equivalent to executing all show class-of-service commands in succession.

show class-of-service classifier

Syntax show class-of-service classifier <name *name*> <type dscp | type exp | type ieee-802.1 | type inet-precedence>

Description For each classifier, display mapping of code point value to forwarding class and loss priority.

Options none—Display all classifiers.

name *name*—(Optional) Display named classifier.

type dscp—(Optional) Display all classifiers of the Differentiated Services Code Point type.

type exp—(Optional) Display all classifiers of the MPLS EXP type.

type ieee-802.1—(Optional) Display all classifiers of the ieee-802.1 type.

type inet-precedence—(Optional) Display all classifiers of the inet-precedence type.

Required Privilege Level view

Output Fields Classifier— Name of the classifier.

Code point type—Type of the classifier. It can be exp, dscp, ieee-802.1, or inet-precedence.

Index—Internal index of the classifier.

Code point—Code point value used for classification.

Forwarding class—The classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Loss priority—Loss priority value used for classification.

Sample Output

```
user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point         Forwarding Class                  Loss priority
  000              best-effort                         low
  001                  best-effort                     high
  010              expedited-forwarding                low
  011              expedited-forwarding                high
  100              assured-forwarding                  low
  101              assured-forwarding                  high
  110              network-control                     low
  111                  network-control                 high

Classifier: users-ieee802.1, Code point type: ieee-802.1
  Code point         Forwarding class                    Loss priority
  100                expedited-forwarding                low
```

show class-of-service code-point-aliases

Syntax show class-of-service code-point-alias < dscp | exp | ieee-802.1 | inet-precedence >

Description Display mapping of code point aliases to corresponding bit patterns.

Options none—Display code point aliases of all code point types.

dscp—(Optional) Display Differentiated Services Code Point aliases.

exp—(Optional) Display MPLS EXP code point aliases.

ieee-802.1—(Optional) Display IEEE-802.1 code point aliases.

inet-precedence—(Optional) Display IPv4 precedence code point aliases.

Required Privilege Level view

Output Fields Code point type—Type of the code points displayed. It can be dscp, exp, ieee-802.1, or inet-precedence.

Alias—Alias for a bit pattern.

Bit pattern—Bit pattern for which the alias is displayed.

Sample Output

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias              Bit pattern
  af11               100
  af12               101
  be                 000
  be1                001
  cs6                110
  cs7                111
  ef                 010
  ef1                011
  nc1                110
  nc2                111
```

show class-of-service drop-profile

Syntax show class-of-service drop-profile < profile-name *profile-name* >

Description Display data points for each Random Early Detection (RED) drop profile.

Options none—Display all drop profiles.

profile-name *profile-name*—(Optional) Display named profile only.

Required Privilege Level view

Output Fields Drop profile—Name of a drop profile.

Type—Type of this drop profile. It can be discrete or interpolated.

Index—Internal index of this drop profile.

Fill level—Percentage fullness of a queue.

Drop probability—Drop probability at this fill level.

Sample Output

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
        100                 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
          0                   0
          1                   1
          2                   2
          4                   4
          5                   5
          6                   6
          8                   8
         10                  10
         12                  15
         14                  20
         15                  23
         ... 64 entries total
         90                  96
         92                  96
         94                  97
         95                  98
         96                  98
         98                  99
         99                  99
        100                 100
```

show class-of-service forwarding-class

Syntax show class-of-service forwarding-class

Description Display mapping of forwarding class name to queue number.

Required Privilege Level view

Output Fields Forwarding class—The classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Queue—Queue corresponding to the forwarding class name.

Sample Output

```
user@host> show class-of-service forwarding-class
Forwarding class                    Queue
  my-forwarding-class                   0
  his-forwarding-class                  1
  her-forwarding-class                  2
  the-forwarding-class                  3
```

show class-of-service forwarding-table

Syntax show class-of-service forwarding-table

Description Display entire class-of-service configuration as it exists in the forwarding table. Executing this command is equivalent to executing all show class-of-service forwarding-table commands in succession.

Required Privilege Level view

Output Fields and Sample Output

show class-of-service forwarding-table classifier on page 493
show class-of-service forwarding-table classifier mapping on page 494
show class-of-service forwarding-table drop-profile on page 495
show class-of-service forwarding-table rewrite-rule on page 496
show class-of-service forwarding-table rewrite-rule mapping on page 497
show class-of-service forwarding-table scheduler-map on page 498

show class-of-service forwarding-table classifier

Syntax show class-of-service forwarding-table classifier

Description Display mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.

Required Privilege Level view

Output Fields Classifier table index—Index of the classifier table.

Entries—Total number of entries.

Table type—Type of code points in the table. It can be DSCP, EXP, IEEE 802.1, or IPv4 precedence.

Entry #—Entry number.

Code point—Code point value used for classification.

Queue #—The queue in which the code point is assigned.

PLP—Packet loss priority value used for classification.

Sample Output

```
user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP

Entry #   Code point   Queue #   PLP
   0        000000       0        0
   1        000001       0        0
   2        000010       0        0
   3        000011       0        0
   4        000100       0        0
   5        000101       0        0
   6        000110       0        0
   7        000111       0        0
   8        001000       0        0
   9        001001       0        0
  10        001010       1        1
  11        001011       0        0
  ...
  60        111100       0        0
  61        111101       0        0
  62        111110       0        0
  63        111111       0        0
```

show class-of-service forwarding-table classifier mapping

Syntax show class-of-service forwarding-table classifier mapping

Description For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Options none

Required Privilege Level view

Output Fields Table index/—If type is Fixed, this is the queue number to which the interface is mapped; for all other types, this is the classifier index number.

Interface—Name of logical interface.

Index—Logical interface index.

Q num—Queue number to which this entry is assigned.

Table type—Type of code points in the table. It can be DSCP, EXP, IEEE 802.1, or IPv4 precedence.

Sample Output

```
user@host> show class-of-service forwarding-table classifier mapping
                                Table index/
Interface      Index     Q num      Table type
so-5/0/0.0        10     62436      DSCP
so-0/1/0.0        11     62436      DSCP
so-0/2/0.0        12         1      Fixed
so-0/2/1.0        13     62436      DSCP
so-0/2/1.0        13     62437      IEEE 802.1
so-0/2/2.0        14     62436      DSCP
so-0/2/2.0        14     62438      IPv4 precedence
```

show class-of-service forwarding-table drop-profile

Syntax show class-of-service forwarding-table drop-profile

Description Display the data points of all Random Early Detection (RED) drop profiles as they exist in the forwarding table.

Required Privilege Level view

Output Fields RED drop profile index—Index of this drop profile.

entries—Number of entries in a particular Random Early Detection (RED) drop profile index.

Entry—Drop profile entry number.

Fullness(%)—Percentage fullness of a queue.

Drop Probability(%)—Drop probability at this fill level.

Sample Output

```
user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
                              Drop
Entry      Fullness(%)   Probability(%)
    0              100              100

RED drop profile index: 8742, # entries: 3
                              Drop
Entry      Fullness(%)   Probability(%)
    0               10               10
    1               20               20
    2               30               30

RED drop profile index: 24627, # entries: 64
                              Drop
Entry      Fullness(%)   Probability(%)
    0                0                0
    1                1                1
    2                2                2
    3                4                4
   ...
   61               98               99
   62               99               99
   63              100              100

RED drop profile index: 25393, # entries: 64
                              Drop
Entry      Fullness(%)   Probability(%)
    0                0                0
    1                1                1
    2                2                2
    3                4                4
   ...
   61               98               98
   62               99               99
   63              100              100
```

show class-of-service forwarding-table rewrite-rule

Syntax show class-of-service forwarding-table rewrite-rule

Description Display mapping of queue number and loss priority to code point value for each rewrite rule as exists in the forwarding table.

Required Privilege Level view

Output Fields Rewrite table index—Index for this rewrite rule.

entries—Number of entries in this rewrite rule.

Table type—Type of code points in the table. It can be DSCP, EXP, IEEE 802.1, or IPv4 precedence.

Q#—Queue number this entry is assigned to.

Low bits—Code point value for low priority loss profile.

State—State of this code point. It can be either enabled (rewritten) or disabled.

High bits—Code point value for high priority loss profile.

Sample Output

```
user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q#     Low bits  State      High bits  State
 0      000111  Enabled      001010   Enabled
 2      000000  Disabled     001100   Enabled
 1      101110  Enabled      110111   Enabled
 3      110000  Enabled      111000   Enabled
```

show class-of-service forwarding-table rewrite-rule mapping

Syntax show class-of-service forwarding-table rewrite-rule mapping

Description For each logical interface, display the table identifier of the rewrite rule mapping for each given code point type.

Required Privilege Level view

Output Fields Interface—Name of logical interface.

Index—Logical interface index.

Table index—Rewrite table index.

Type—Type of classifier. It can be unknown, DSCP, EXP, IEEE 802.1, IPv4 precedence, Fixed, or illegal.

Sample Output

```
user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface      Index   Table index  Type
so-5/0/0.0        10       3753  DSCP
so-0/1/0.0        11       3753  DSCP
so-0/2/0.0        12       3753  DSCP
so-0/2/1.0        13       3753  DSCP
so-0/2/2.0        14       3753  DSCP
so-0/2/3.0        15       3753  DSCP
```

show class-of-service forwarding-table scheduler-map

Syntax show class-of-service forwarding-table scheduler-map

Description For each physical interface, display the scheduler map information as it is in the forwarding table.

Options none

Required Privilege Level view

Output Fields Interface—Name of the physical interface.

Index—Physical interface index.

Map index—Scheduler map index.

Num of queues—Number of queues defined in this scheduler map.

Entry—Number of this entry in the scheduler map.

Scheduler index—Scheduler policy index.

Queue #—Which queue number this entry is applied to.

Tx rate—Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder, which indicates that the scheduler should receive the remaining bandwidth of the interface.

Max buffer delay—Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or by the keyword remainder, which indicates that the buffer should be sized according to what remains after other scheduler buffer allocations.

High priority is set—If this line appears in the output, the queue priority is high; otherwise it is low.

PLP high—Drop profile index for high priority packet loss profile.

PLP low—Drop profile index for low priority packet loss profile.

TCP PLP high—Drop profile index for high priority TCP packet loss profile.

TCP PLP low—Drop profile index for low priority TCP packet loss profile.

Policy is exact—If this line appears in the output, exact rate limiting is enabled; otherwise no rate limiting is enabled.

Sample Output

```
user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Queue #: 0):
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    High priority is set
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low:8742
    Policy is exact
  Entry 1 (Scheduler index: 38372, Queue #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    High priority is set
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Queue #: 0):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    High priority is set
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
  Entry 1 (Scheduler index: 38372, Queue #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    High priority is set
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
```

show class-of-service interface

Syntax show class-of-service interface *type-fpc/pic/port*

Description Display the logical and physical interface associations for the classifier, rewrite rules and scheduler map objects.

Options none—Display CoS associations for physical and logical interfaces.

type-fpc/pic/port—(Optional) Display CoS associations for named physical and logical interfaces.

Required Privilege Level view

Output Fields Physical interface—Name of a physical interface.

Scheduler map—Name of the scheduler map associated with this interface.

Index—Index of this interface or the internal index of this object.

Logical interface—Name of a logical interface.

Object —Category of an object, one of classifier or rewrite.

Name—Name of an object.

Type—Type of an object. It can be dscp, exp, ieee-802.1. or inet-precedence.

Sample Output

```
user@host> show class-of-service interface so-0/2/3
 Physical interface: so-0/2/3, c: 32
  Scheduler map: my-scheduler-map, Index: 17638

  Logical interface: so-0/2/3.0, Index: 15
   Object         Name                   Type                Index
   Rewrite        my-dscp-rewrite        dscp                 3753
   Classifier     my-dscp-classifier     dscp                62436

user@host> show class-of-service interface so-0/2/3.0
  Logical interface: so-0/2/3.0, Index: 15
   Object         Name                   Type                Index
   Rewrite        my-dscp-rewrite        dscp                 3753
   Classifier     my-dscp-classifier     dscp                 62436
```

show class-of-service rewrite-rule

Syntax show class-of-service rewrite-rule < name *name* > < type dscp | type exp | type inet-precedence >

Description Display mapping of forwarding class and loss priority to code point value.

Options none—Display all rewrite rules.

name *name*—(Optional) Display named rewrite rule.

type dscp—(Optional) Display named rewrite rule of type dscp.

type exp—(Optional) Display named rewrite rule of type exp.

type inet-precedence—(Optional) Display named rewrite rule of type inet-precedence.

Required Privilege Level view

Output Fields Rewrite rule—Name of a rewrite rule.

Code point type—Type of this rewrite rule. This can be exp, dscp or inet-precedence.

Forwarding class—The classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Index—Internal index for this particular rewrite rule.

Loss priority—Loss priority for rewriting.

Code point—Code point value to rewrite.

Rewrite rule—Name of a rewrite rule.

Sample Output

```
user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class                    Loss priority       Code point
  gold                                high                000000
  silver                              low                 110000
  silver                              high                111000
  bronze                              low                 001010
  bronze                              high                001100
  lead                                high                101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
Forwarding class                    Loss priority       Code point
  gold                                low                 000111
  gold                                high                001010
  silver                              low                 110000
  silver                              high                111000
  bronze                              high                001100
  lead                                low                 101110
  lead                                high                110111
```

show class-of-service scheduler-map

Syntax show class-of-service scheduler-map <*name*>

Description Display mapping of scheduler to forwarding class and summary of scheduler parameters for each entry.

Options none—Display all scheduler maps.

name—Display summary of scheduler parameters for each forwarding class the named scheduler is assigned.

Required Privilege Level view

Output Fields Scheduler map—Name of the scheduler map.

Index—Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.

Scheduler—Name of the scheduler.

Forwarding class—The classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Transmit Rate—Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder, which indicates that the scheduler should receive the remaining bandwidth of the interface.

Rate Limit—Rate limiting configuration of the queue. Possible values are none, meaning no rate limiting, and exact, meaning the queue will only transmit at the configured rate.

Maximum buffer delay—The amount of transmit delay (in milliseconds) or buffer size of the queue, as a percentage of the total interface buffer allocation or by the keyword "remainder." This indicates that the buffer should be sized according to what remains after other scheduler buffer allocations.

Priority—Scheduling priority; possible values include low and high.

Drop profiles—This table shows the assignment of drop profile by name and index to a given loss priority and protocol pair.

Loss priority—Packet loss priority for drop profile assignment.

Protocol—Transport protocol for drop profile assignment.

Name—Name of the drop profile.

Sample Output

```
user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

  Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
   Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile

  Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
   Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
   Priority: high
   Drop profiles:
    Loss priority   Protocol    Index    Name
     Low            non-TCP     8724     aa-drop-profile
     Low            TCP         9874     bb-drop-profile
     High           non-TCP     8833     cc-drop-profile
     High           TCP         8484     dd-drop-profile
```

show interfaces extensive (CoS related attributes only)

Syntax show interfaces extensive <interface-name *interface-name*>

Description Display CoS statistics for physical interfaces.

Options none—Display CoS queue statistics for all physical interfaces.

interface-name *interface-name*—(Optional) Name of physical interface to be displayed.

Required Privilege Level view

Output Fields (related to CoS)

Physical interface—Name of the physical interface.

Queue counters—Queue number and its associated forwarding class designation.

Queued packets—Number of packets queued to the given queue.

Transmitted packets—Number of packets transmitted by the given queue.

RED dropped packets—Number of RED (Random Early Detection) dropped packets.

PFE configuration—Information about how the Packet Forwarding Engine is configured.

Destination slot—Slot the card is plugged into.

CoS transmit queue—CoS transmit queue number and its associated forwarding class designation.

Bandwidth%—Configured bandwidth as percentage of the specified queue.

Bandwidth bps—Configured bandwidth in bps (bits per second).

Buffer%—Configured buffer percentage of the specified queue.

Buffer bytes—Configured size of the queue in bytes.

Priority—The queue priority. Possible values are low, high, none or exact. None indicates no rate limiting and exact indicates the queue will only transmit at the configured rate.

Limit—Rate limiting configuration of the queue. Possible values are none, meaning no rate limiting, and exact, meaning the queue will only transmit at the configured rate.

Dropped packets—Number of packets dropped by the specified queue due to Random Early Detection (RED) or tail drop.

Sample Output (related to CoS)

```
user@host> show interfaces extensive ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Up
...
 Queue counters:       Queued packets  Transmitted packets      Dropped packets
   0 best-effort                  103                  103                    0
   1 expedited-fo           616679163            588733189                    0
   2 assured-forw          2648825387           2528227577            120446880
   3 network-cont            27284126             26134172              1149954
...
  PFE configuration:
    Destination slot: 0, PLP byte: 1 (0x00)
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                                 %          bps   %        bytes
    0 best-effort               95            0  95            0      low    none
    1 expedited-forwarding       0            0   0            0      low    none
    2 assured-forwarding         0            0   0            0      low    none
    3 network-control            5            0   5            0      low    none
```

show interfaces queue

Syntax show interfaces queue < interface-name *interface-name*><forwarding-class *forwarding-class*>

Description Display CoS information per physical interface.

Options none—Shows detailed CoS queue statistics for all physical interfaces.

interface-name *interface-name*—(Optional) Shows detailed CoS queue statistics for named physical interface.

forwarding-class *forwarding-class*—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for queue associated with named forwarding class.

Required Privilege Level view

Output Fields Physical interface—Name of the physical interface.

Interface index—Physical interface's index number, which reflects its initialization sequence.

SNMP ifIndex—SNMP index number for the interface.

Forwarding class—The classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Queue—Queue number.

Queued Packets—Number of packets queued to this queue.

Queued Bytes—Number of bytes queued to this queue.

Transmitted Packets—Number of packets transmitted by this queue.

Transmitted Bytes—Number of bytes transmitted by this queue.

Tail-dropped Packets—Number of packets dropped due to tail drop.

RED-dropped packets—Number of packets dropped due to Random Early Detection (RED).

- Low, non-TCP—Number of low loss priority, non-TCP packets dropped due to RED.
- Low, TCP—Number of low loss priority, TCP packets dropped due to RED.
- High, non-TCP—Number of high loss priority, non-TCP packets dropped due to RED.
- High, TCP—Number of high loss priority, TCP packets dropped due to RED.

RED-dropped bytes—Number of bytes dropped due to RED.

- Low, non-TCP—Number of low loss priority, non-TCP bytes dropped due to RED.
- Low, TCP—Number of low loss priority, TCP bytes dropped due to RED.
- High, non-TCP—Number of high loss priority, non-TCP bytes dropped due to RED.
- High, TCP—Number of high loss priority, TCP bytes dropped due to RED.

Sample Output

```
user@host> show interfaces queue so-0/2/3 forwarding-class expedited-forwarding
Physical interface: so-0/2/3, Enabled, Physical link is Up
  Interface index: 32, SNMP ifIndex: 60
Forwarding Class: expedited-forwarding, Queue: 1
  Queued:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
  Transmitted:
    Packets              :                    0                    0 pps
    Bytes                :                    0                    0 bps
    Tail-dropped packets :                    0                    0 pps
    RED-dropped packets  :                    0                    0 pps
     Low, non-TCP        :                    0                    0 pps
     Low, TCP            :                    0                    0 pps
     High, non-TCP       :                    0                    0 pps
     High, TCP           :                    0                    0 pps
    RED-dropped bytes    :                    0                    0 bps
     Low, non-TCP        :                    0                    0 bps
     Low, TCP            :                    0                    0 bps
     High, non-TCP       :                    0                    0 bps
     High, TCP           :                    0                    0 bps
```

Part 6
Protocols

- Protocol-Independent Routing Monitoring and Troubleshooting on page 511
- IS-IS Monitoring and Troubleshooting on page 573
- OSPF Monitoring and Troubleshooting on page 593
- BGP Monitoring and Troubleshooting on page 613
- RIP Monitoring and Troubleshooting on page 625
- RIPng Monitoring and Troubleshooting on page 629
- IPv6 Monitoring and Troubleshooting on page 635
- IP Multicast Monitoring and Troubleshooting on page 639

Chapter 34

Protocol-Independent Routing Monitoring and Troubleshooting

Table 65 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot protocol-independent routing properties. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

> **Note**
> In Table 65, the variants of the show route commands listed in the category "Routing Table Information" are all filters that select the specific information that is displayed from the routing tables. The other show route commands (specifically, show route forwarding-table and show route martians) are independent commands and cannot be used in conjunction with any of the filter options.

Table 65: Commands for Monitoring Protocol-Independent Routing Properties

Task Category	Task or Information to Monitor	Command
Routing Table Information	Information about the entries in the routing tables.	show route on page 516
	Routes transmitted by a particular routing protocol.	show route advertising-protocol on page 527
	Routes containing a specified AS path.	show route aspath-regex on page 529
	Best route to the specified address or range of addresses.	show route best on page 531
	Routes containing members of a specified BGP community.	show route community on page 532
	Routes that have been damped.	show route damping on page 534
	Routes that exactly match the specified address or range of addresses.	show route exact on page 537
	Display list of instances or routing tables that are importers or exporters of routes.	show route export on page 539
	Display target communities for which auto-export is currently distributing routes.	show route export vrf-target on page 541
	Routes that are currently inactive.	show route inactive on page 549
	Routing instance information.	show route instance on page 550
	Routes that form a label-switched path.	show route label-switched-path on page 552
	Routes that contain the specified next hop.	show route next-hop on page 555
	Routes exiting the router through the specified interface.	show route output on page 556
	Routes learned by the specified protocol.	show route protocol on page 557
	Routes in a range of destination prefixes.	show route range on page 559
	Routes received by a particular routing protocol.	show route receive-protocol on page 561
	Entries in the next-hop resolution database.	show route resolution on page 563
	Routes learned from the specified source.	show route source-gateway on page 565
	Statistics about the routes in all routing tables.	show route summary on page 566
	Routes in a particular routing table.	show route table on page 567
	High-level summary of routing table information.	show route terse on page 571
Forwarding Table Information	Information about the entries in the kernel's forwarding table.	show route forwarding-table on page 543
	Clear a route entry from the kernel's forwarding table.	clear route forwarding-table on page 513
Martian Information	Information about martian addresses.	show route martians on page 554
AS Paths	Known AS paths.	show as-path on page 514

clear route forwarding-table

Syntax clear route forwarding-table *<destination-prefix > <nexthop>*

Description Delete a route entry from the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine.

Options *destination-prefix*—(Optional) IP address of the destination to delete.

nexthop—(Optional) Next hop for the route.

Required Privilege Level maintenance

Sample Output

```
user@host> show route forwarding-table
Internet:
Destination             Type RtRef Nexthop           Type Index NhRef Netif
...
123.456.0.0/18          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.457.0.0/18          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.123.0.0/17          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.123.24.0/21         user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.125.0.0/16          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.129.0.0/18          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.0.0/17          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.24.0/22         user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.128.0/18        user     0 111.222.5.254     ucst    19 55371 fxp0.0
...
user@host> clear route forwarding-table 123.456.0.0
user@host> show route forwarding-table
Internet:
Destination             Type RtRef Nexthop           Type Index NhRef Netif
...
123.457.0.0/18          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.123.0.0/17          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.123.24.0/21         user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.125.0.0/16          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.129.0.0/18          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.0.0/17          user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.24.0/22         user     0 111.222.5.254     ucst    19 55371 fxp0.0
123.130.128.0/18        user     0 111.222.5.254     ucst    19 55371 fxp0.0
...
```

show as-path

Syntax show as-path

Description Display the distribution of AS paths that the local routing is using (usually through the routing table). Use this command when debugging refcount problems for AS paths and for seeing how AS paths have been manipulated through policy (through the as-path-prepend action) or aggregation.

Required Privilege Level view

Output Fields Total AS paths—Number of AS paths.

Path—AS path.

<flags>—Information about the AS path:

- ASLoop—Path contains an AS loop.
- Atomic—Path includes the ATOMIC_AGGREGATE path attribute.
- Local—Path was created by local aggregation.

Refs—Path reference count.

ASes—Number of AS paths that the router knows.

Segments—Length of the AS segment descriptor.

Overhead—Efficiency of AS storage.

Sample Output

```
user@host> show as-path
Total AS paths: 24744
Path: 234 234 3561 6503 278 3596 I
        Refs: 2 ASes: 6 Segments: 1 Overhead: 64
Path: 267 234 3561 4926 4926 4926 10834 5648 I <ASLoop>
        Refs: 2 ASes: 8 Segments: 1 Overhead: 64
Path: 234 234 3561 4926 4926 4926 10834 5648 I <ASLoop>
        Refs: 2 ASes: 8 Segments: 1 Overhead: 64
Path: 2914 4000 8068 8072 I
        Refs: 3 ASes: 4 Segments: 1 Overhead: 64
Path: 2914 701 6388 3464 I
        Refs: 22 ASes: 4 Segments: 1 Overhead: 64
Path: 267 234 1239 6347 11404 11404 I <ASLoop>
        Refs: 2 ASes: 6 Segments: 1 Overhead: 64
Path: 234 234 1239 6347 11404 11404 I <ASLoop>
        Refs: 2 ASes: 6 Segments: 1 Overhead: 64
Path: 267 234 1239 2516 4675 I <Atomic>
        Aggregator: 4675 202.233.0.5
        Refs: 2 ASes: 5 Segments: 1 Overhead: 64
Path: 267 234 1239 2516 4675 I <Atomic>
        Aggregator: 4675 202.249.2.37
        Refs: 2 ASes: 5 Segments: 1 Overhead: 64
Path: 234 234 1239 2516 4675 I <Atomic>
        Aggregator: 4675 202.249.2.37
        Refs: 2 ASes: 5 Segments: 1 Overhead: 64
Path: 234 234 1239 2516 4675 I <Atomic>
        Aggregator: 4675 202.233.0.5
        Refs: 2 ASes: 5 Segments: 1 Overhead: 64
Path: 267 234 1239 701 10275 ?
        Aggregator: 10275 157.130.192.246
        Refs: 4 ASes: 5 Segments: 1 Overhead: 64
...
```

show route

Syntax show route <brief | detail | extensive> <all | hidden> <*destination-prefix*> <*filters*> <*routing-table-name*>

Description Display the active entries in the routing tables.

Options none—Display all active entries in the routing tables.

all—(Optional) Display information about all routes, including hidden entries.

brief—(Optional) Display brief route information.

detail—(Optional) Display detailed route information.

extensive—(Optional) Display very detailed route information.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

filters—(Optional) One or more filters that select specific information to display from the routing tables. To simplify the readability and explanation of the filters, each filter is treated as an individual command in this manual, even though you can combine filters in a single show route command. The following is the syntax of the show route command including all filters:

show route <*destination-prefix*> <summary | detail | extensive> <best | exact | range> <advertising-protocol *protocol neighbor-address*> <aspath-regex *regular-expression*> <community *as-number*:*community-value*> <damping> <export> <forwarding-table> <inactive> <instance> <label-switched-path> <martians> <next-hop> <output> <protocol *protocol*> <range> <receive-protocol *protocol neighbor-address*> <source-gateway *address*> <summary> <table *routing-table-name*> <terse>

hidden—(Optional) Display hidden route information.

routing-table-name—(Optional) Display information about a particular routing table.

Default: brief

Required Privilege Level view

Sample Output: show route

Sample Output: show route brief on page 523
Sample Output: show route detail, non-route reflector case on page 524
Sample Output: show route extensive, non-route reflector case on page 524
Sample Output: show route detail, route reflector case on page 525
Sample Output: show route extensive, route reflector case on page 525
Sample Output: show route extensive, RSVP link-protection case on page 525
Sample Output: show route extensive on page 526

Options at a Glance Table 66 summarizes which information is included in each show route command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 66: Show Route Output Field Summary

Options	Field Description
Brief	[*protocol/preference*]—Protocol from which the route was learned and the route's preference value.
Detail Extensive	Age—How long the route has been known.
Detail Extensive	Announcement bits—List of protocols that will announce this route.
Detail Extensive	AS path—AS path through which the route was learned.
Detail Extensive	AS path: I <*Originator*>—(For route reflected output only) Originator ID sent by the route reflector.
Detail	BGP next hop—Next hop path attribute of the route.
Detail Extensive	Cluster list—(For router reflected output only) Cluster ID sent by the route reflector.
Detail	Communities—Community path attribute of the route.
Extensive	Destination class—Prefix destination class for class billing.
Detail Extensive	*destination-prefix* (entry, announced)—Route address. entry is the number of routes for this destination, and announced is the number of routes being announced for this destination.
All	destinations—Number of destinations for which there are routes in the routing table.
Detail Extensive	domain-id—A unique, configurable number that identifies the OSPF domain. This is configured using the BGP extended community attribute 0x0105. The default is 0.
Detail Extensive	domain-id-vendor—A unique, configurable number that identifies the OSPF domain. This is configured using extended community attribute 0x8105.
Detail Extensive	Inactive reason—Flags for this route, which was not selected as best for a particular destination.
Extensive	Indirect nexthops—An index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
Detail Extensive	Label operation—MPLS label and operation occurring at this router. The operation can be pop, push, or swap.
Detail	Local AS—AS number of the local router.
Detail	Localpref—Local preference path attribute of the route.
Detail	Metric2—For routes learned from BGP, this is the IGP metric.
Detail	Metric—For routes learned from BGP, this is the MED metric.
Detail Extensive	Next hop type—Type of next hop. It can be discard, Local, Interface, Unusable or Reject.
Detail Extensive	Nexthop—Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.
Detail Extensive	Not Best in its group—After BGP path selection process is run on a group of incoming ASs, one will be chosen as the best in group; the others will be "not best in group."
Detail Extensive	origin—Identifies where the route came from. origin is used for VPNs.
Detail Extensive	Originator ID—(For router reflected output only) Address of router that originally sent route to the route reflector.
Detail	Peer AS—AS number of the peer router.
Detail Extensive	Preference—Preference value of the route.
Detail Extensive	protocol nexthop—Network layer address of a remote router that advertised the prefix. This address is used to recurse through and derive a forwarding next hop.
Detail Extensive	Route Distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.
Detail	Router ID—BGP router ID as advertised by the neighbor in the open message.
Detail Extensive	route-type-vendor—Displays the area number, OSPF route type and option of the route. This is configured using the BGP extended community attribute 0x8000.
All	routes—Number of routes in the routing table.
All	*routing-table-name*—Name of the routing table; for example, inet.0.

Options	Field Description
Detail Extensive	rte-type—Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is *area-number:ospf-route-type:options*.
Detail Extensive	Source—Source address of the route.
Detail Extensive	State—Flags for this route.
Detail Extensive	Static | Direct | Local | *protocol-name*—How the route was learned.
Detail Extensive	target—Defines which VPN the route participates. The format is *area-number:ospf-route-type:options*.
Detail Extensive	Task—Name of the task that owns the route.
Brief	*time*—How long the route has been known.
Extensive	TSI—Header for protocol-specific information.
Detail Extensive	unknown IANA—Incoming IANA codes with a value between 0x1-0x7fff cannot be identified. This code of the BGP extended community attribute is accepted, but it is not recognized.
Detail Extensive	unknown OSPF vendor community—Incoming IANA codes with a value above 0x8000 cannot be identified.This code of the BGP extended comminute attribute is accepted, but it is not recognized.
Detail Extensive	Update source—The last tiebreaker is the lowest IP address value.
Detail Extensive	VPN-Label—VPN label learned in the advertisement.
Detail Extensive	Weight—Measure of the importance of the next hop. The lower the value, the more importance.

Output Fields

routing-table-name—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

Destination class—(Extensive output only) Prefix destination class for class billing.

[*protocol*/*preference*]—(Brief output only) Protocol from which the route was learned and the route's preference value.

time—(Brief output only) How long the route has been known.

destination-prefix (entry, announced)—(Detail and Extensive output only) Route address. The value of entry is the number of routes for this destination, and announced is the number of routes being announced for this destination.

Static | Direct | Local | *protocol-name*—(Detail and Extensive output only) How the route was learned.

Preference—(Detail and Extensive output only) Preference value of the route.

Route distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.

Weight—(Detail and Extensive output only) Measure of the importance of the next hop. The lower the value, the more importance.

- 1-9,999—For regular routes
- 10,000-19,999—For MPLS fast reroute routes
- 20,000-29,999—For link protection routes

Nexthop—(Detail and Extensive output only) Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Next hop type—(Detail and Extensive output only) Type of the next hop. It can be Discard, Local, Interface, Unusable or Reject.

Indirect nexthops—(Extensive output only) An index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.This information is present only if the path in the routing table includes indirect next hop information.

Source—(Detail and Extensive output only) Source address of the route.

Label operation—(Detail and Extensive output only) MPLS label and operation occurring at this router. The operation can be pop, push, or swap.

Protocol nexthop—Network layer address of a remote router that advertised the prefix. This address is used to recurse through and derive a forwarding next-hop.

State—(Detail and Extensive output only) Flags for this route. It can be one or more of the following:

- Accounting—Route needs accounting.
- Active—Route is active.
- Clone—Route is a clone.
- Delete—Route deleted.
- Ex—Exterior route.
- Hidden—Route not used because of routing policy.
- IfCheck—Route needs forwarding RPF check.
- Initial—Route being added.
- Int—Interior route.

- Martian—Route is a martian.
- MartianOK—Route exempt from martian filtering.
- NoReadvrt—Route not to be advertised.
- NotBest—Route not chosen because it does not have the lowest MED metric.
- NotInstall—Route not to be installed in the forwarding table.
- Pending—Route pending because of holddown on another route.
- Release—Route scheduled for release.
- Secondary—Route not a primary route.

Inactive reason—(Detail and Extensive output only) Flags for this route, which was not selected as best for a particular destination. It can be one or more of the following:

- Always Compare MED—Always compare MED is enabled and a path with lower MED is available.
- AS path—Shorter AS path available.
- Cisco Non-deterministic MED selection—Cisco non-deterministic MED is enabled, and a path with lower MED is available.
- Cluster list length—(Route reflection only) Length of cluster list sent by the route reflector.
- IGP metric—Path through next hop with lower IGP metric available.
- Interior > Exterior > Exterior via Interior—A direct, static, IGP, or EBGP path is available.
- Local Preference—Path with a higher local preference value available.
- Nexthop address—Path with lower metric next hop available.
- No difference—Path from neighbor with lower IP address available.
- Not Best in its group—After BGP path selection process is run on a group of incoming ASs, one will be chosen as the best in group; the others will be “not best in group.”
- Number of gateways—Path with greater number of next hops available.
- Origin—Path with lower origin code available.
- RIB preference—Route from a higher numbered routing table available.
- Route Distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.
- Route Metric or MED comparison—Route with lower metric or MED available.
- Route Preference—Route with lower preference value available.

- Router ID—Path through neighbor with lower ID available.
- Update source—The last tiebreaker is the lowest IP address value.
- Unusable path—Path not usable.

Local AS—(Detail output only) AS number of the local router.

Peer AS—(Detail output only) AS number of the peer router.

Age—(Detail and Extensive output only) How long the route has been known.

Metric—(Detail output only) For routes learned from BGP, this is the MED metric.

Metric2—(Detail output only) For routes learned from BGP, this is the IGP metric.

Task—Name of the protocol that has added the route.

Announcement bits—(Detail and Extensive output only) List of protocols that announce this route. *n*-Resolve inet.*x* indicates that the route is used for route resolution for next hops found in the routing table inet.*x*. *n* is an index used by Juniper Networks Customer Support only.

AS path—(Detail and Extensive output only) AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path was originated:

- I—IGP.
- E—EGP.
- ?—Incomplete; typically, the AS path was aggregated.

Communities—(Detail output only) Community path attribute of the route. See Internet draft draft-rosen-vpns-ospf-bgp-mpls-nn.txt for more information on these extended community attributes.

- target—Defines which VPN the route participates in.target has the format *32-bit IP address:16-bit number*. For example, 10.19.0.0:100.
- origin—Identifies where the route came from. origin is used for VPNs.
- rte-type—Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is *area-number:ospf-route-type:options*.
  - *area-number*—4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A non-zero value identifies the route as being internal to the OSPF domain, and as being within the identified area. Area numbers are relative to a particular OSPF domain.
    - *ospf-route-type*—1 byte, encoded as follows:
    - 1 or 2—Intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA)
    - 3—Summary routes

- 5—External routes (area number must be 0)
- 7—NSSA routes
- 129—Sham Link Endpoint Addresses

- *options*—1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.

- route-type-vendor—Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000.The format is *area-number:ospf-route-type:options*.

  - *area-number*—4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A non-zero value identifies the route as being internal to the OSPF domain, and as being within the identified area. Area numbers are relative to a particular OSPF domain.

    - *ospf-route-type*—1 byte, encoded as follows:
    - 1 or 2—Intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA)
    - 3—Summary routes
    - 5—External routes (area number must be 0)
    - 7—NSSA routes
    - 129—Sham Link Endpoint Addresses

  - *options*—1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.

- domain-id—A unique, configurable number that identifies the OSPF domain.This is configured using the BGP extended community attribute 0x0x0105. domain-id is present unless the VRF has the (default) domain identifier value of zero. In this case, the attribute can be omitted. This attribute is encoded with a two-byte type field, and its type is either 0x0005, 0x0105, or 0x0205. The value of the domain identifier is carried in the “local administrator” subfield. The type 0x8005 should be accepted as well, to ensure backward compatibility, and treated as if it were 0x0005.

- domain-id-vendor—A unique, configurable number that identifies the OSPF domain. This is configured using extended community attribute 0x8105.

- unknown IANA—Incoming IANA codes with a value between 0x1-0x7fff cannot be identified. This code of the BGP extended community attribute is accepted, but it is not recognized.

- unknown OSPF vendor community—Incoming IANA codes with a value above 0x8000 cannot be identified.This code of the BGP extended community attribute is accepted, but it is not recognized.

VPN-Label—VPN label learned in the advertisement.

Localpref—(Detail output only) Local preference path attribute of the route.

Route ID—(Detail output only) BGP router ID as advertised by the neighbor in the open message.

AS path: I <Originator>—(For route reflected output only; Detail and Extensive output only) Originator ID sent by the route reflector.

Cluster list—(For router reflected output only; Detail and Extensive output only) Cluster ID sent by the route reflector.

Originator ID—(For router reflected output only; Detail and Extensive output only) Address of router that originally sent route to the route reflector.

TSI—(Extensive output only) Header for protocol-specific information.

Sample Output: show route brief

```
user@host> show route brief
inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:00:08
                    > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:00:09
                    > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:00:09
                      Local
12.12.12.21/32     *[Local/0] 00:00:08
                      Reject
13.13.13.13/32     *[Direct/0] 00:00:09
                    > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:00:09
                      Local
13.13.13.21/32     *[Local/0] 00:00:09
Local
13.13.13.22/32     *[Direct/0] 00:00:09
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:00:09
                    > via lo0.0
111.222.5.0/24     *[Direct/0] 00:00:09
                    > via fxp0.0
111.222.5.59/32    *[Kernel/-2] 00:00:09
                    > to 111.222.5.62 via fxp0.0
111.222.5.62/32    *[Kernel/-2] 00:00:09
                    > to 111.222.5.62 via fxp0.0
111.222.5.81/32    *[Local/0] 00:00:09
                      Local
224.0.0.5/32       *[OSPF/10] 00:00:09, metric 1

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.5081.00/160
                   *[Direct/0] 00:00:09
                    > via lo0.0
```

Sample Output: show route detail, non-route reflector case

```
user@host> show route detail
1.0.0.0/8 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Source: 207.17.136.192
                Nexthop: 192.168.4.254 via 192.168.5.0, selected
                Protocol Nexthop: 207.17.136.192 Indirect nexthop: 851b110 47
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 2:32       Metric: 0       Metric2: 0
                Task: BGP_10458.207.17.136.192+179
                Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
                AS path: 3944 7777 I
                Communities: 7777:7777
                Localpref: 100
                Router ID: 10.255.245.88
```

Sample Output: show route extensive, non-route reflector case

```
user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

1.0.0.0/8 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(47)}
Page 0 idx 1 Type 1 val a4d01c0
    Nexthop: 207.17.136.192
    MED: 0
    Localpref: 100
    AS path: 3944 7777 I
    Communities: 7777:7777
    Cluster ID: 1.1.1.1
    Originator ID: 10.255.245.88
Path 1.0.0.0 from 207.17.136.192 Vector len 4.  Val: 1

        *BGP    Preference: 170/-101
                Source: 207.17.136.192
                Nexthop: 192.168.4.254 via so-0/1/0.0, weight 1, selected
                Protocol Nexthop: 207.17.136.192 Indirect nexthop: 851b110 47
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 3:46       Metric: 0       Metric2: 0
                Task: BGP_10458.207.17.136.192+179
                Announcement bits (3): 0-KRT 3-BGP.0.0.0.0+179 4-Resolve inet.0
                AS path: 3944 7777 I
                Communities: 7777:7777
                Localpref: 100
                Router ID: 10.255.245.88
                Indirect nexthops: 1
                        Protocol nexthop: 207.17.136.192 Metric: 0 Indirect
                          nexthop: 851b110 47
                        Indirect path forwarding nexthops: 1
                                Nexthop: 192.168.4.254 via fxp0.0
```

Sample Output: show route detail, route reflector case

```
user@host> show route 16/8 detail
1.0.0.0/8 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol Nexthop: 207.17.136.192 Indirect nexthop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 48 Metric: 0       Metric2: 0
                Task: BGP_10458.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 3944 7777 I <Originator>
                Cluster list:  1.1.1.1
                Originator ID: 10.255.245.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 4.4.4.4
```

Sample Output: show route extensive, route reflector case

```
user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}

        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol Nexthop: 207.17.136.192 Indirect nexthop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_10458.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 3944 7777 I <Originator>
                Cluster list:  1.1.1.1
                Originator ID: 10.255.245.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 4.4.4.4
                Indirect nexthops: 1
                        Protocol nexthop: 207.17.136.192 Metric: 0 Indirect
                        nexthop: 84ac908 40
                        Indirect path forwarding nexthops: 0
                                Next hop type: Discard
```

Sample Output: show route extensive, RSVP link-protection case

```
user@host> show route extensive
100002 (1 entry, 1 announced)
TSI:
KRT in-kernel 100002 /36 -> {0.0.0.0}
         *RSVP   Preference: 7
                 Nexthop: via so-0/1/2.0, weight 1, selected
                 Label-switched-path Bypass_to_192.168.207.225
                 Label operation: Pop 100001
                 State: <Active Int>
                 Age: 27:16      Metric: 1
                 Task: RSVP
                 Announcement bits (1): 0-KRT
                 AS path: I
```

Sample Output: show route extensive

```
user@host> show route vpn02 extensive
0.0.0.0/0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0.0.0.0/0 -> {}
        *Static Preference: 5
                Next hop type: Discard
                State: <Active NoReadvrt Int Ext>
                Age: 1:15:11    Metric: 0
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

CE_A.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

10.255.14.172/32 (1 entry, 1 announced)
TSI:
OSPF LSA ID 10.255.14.172
KRT in-kernel 10.255.14.172/32 -> {indirect(103)}
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.14.176:1
                Route Distinguisher: 10.255.14.176:1
                Source: 10.255.14.176
                Next hop: via t3-0/0/1.0,, selected
                Label operation: Push 100008, Push 100038(top)
                Protocol next hop: 10.255.14.176
                Push 100008
                 Indirect next hop: 84162a8 103
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 11:26      Metric: 2       Metric2: 1
                Task: BGP_69.10.255.14.176+179
                Announcement bits (2): 0-CE_A-OSPF 3-KRT
                AS path: I
                Communities: target:10.19.0.0:100 route-type:0.0.0.0:1:0
                    domain-id-vendor:1.2.3.4
                VPN Label: 100008
                Localpref: 100
                Router ID: 10.255.14.176
```

show route advertising-protocol

Syntax show route advertising-protocol *protocol neighbor-address* <detail>

Description Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol. The information displayed reflects the routes that the routing table has exported into the routing protocol and that were filtered by that protocol's export routing policy statements. (In the figure "Importing and Exporting Routing Policies" in the section "How to Construct Routing Policy" in the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*, these are the routes shown in the right "Protocol" box.)

Options detail—(Optional) Display additional information that was advertised to the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were advertised, including communities, route reflection, and aggregator.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route. It can be bgp, dvmrp, msdp, pim, or rip.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Nexthop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

Route-distinguisher—The unique route distinguisher associated with this routing instance.

VPN Label—Name of label assigned to this VPN.

Communities—Community path attribute of the route.

Sample Output

```
user@host> show route advertising-protocol bgp 111.222.1.3
inet.0: 46498 destinations, 46498 routes (46496 active, 0 holddown, 2 hidden)
Prefix             Next hop           MED        Lclpref    AS path
15.0.0.1/32        111.222.1.1                        69 IGP
```

Sample Output

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail

bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
111.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     Nexthop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
111.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     Nexthop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
```

show route aspath-regex

Syntax

show route aspath-regex *regular-expression*
show route aspath-regex "*regular-expression* ..."

Description

Display the entries in the routing table that match the specified BGP AS path regular expression.

Options

regular-expression—Regular expression that matches an entire AS path.

You can specify an AS path term in one of the following ways:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table "AS Path Regular Expression Operators" in the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*. The following list summarizes these operators:

- {*m*,*n*}—At least *m* and at most *n* repetitions of the AS path term.
- {m}—Exactly *m* repetitions of the AS path term.
- {m,}—*m* or more repetitions of the AS path term.
- *—Zero or more repetitions of an AS path term.
- +—One or more repetitions of an AS path term.
- ?—Zero or one repetition of an AS path term.
- *aspath_term* | *aspath_term*—Match one of the two AS path terms.

When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

```
show route aspath-regex ".* 234 .*"
```

Required Privilege Level

view

Output Fields

routing-table-name—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

time—How long the route has been known.

localpref—Local preference value included in the route.

from—Address and interface from which route was learned.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

AS Path—AS path included in the route.

[ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the router.

{ }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.

( )—Parentheses enclose a confederation.

([ ])—Enclose a confederation set.

Sample Output

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25     *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
                    [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

```
user@host> show route aspath-regex ".* 234 3561 .*"

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.88.0.0/18       *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 7725 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.92.0.0/19       *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 IGP
...
```

show route best

Syntax show route best *destination-prefix*

Description Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options *destination-prefix*—Address or range of addresses.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

time—How long the route has been known.

Sample Output

```
user@host> show route best 111.222/24

inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:09:41
                    > to 111.222.5.254 via so-2/0/1
```

show route community

Syntax

show route community *as-number:community-value*
show route community "*as-number:community-value*"

Description

Display the route entries in each routing table that are members of a BGP community. The Communities: field in the show route detail command output lists the communities that the route is a member of, if any.

Options

as-number:community-value—One or more community identifier. *as-number* is the AS number, and *community-value* is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks.

Required Privilege Level

view

Output Fields

routing-table-name—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

time—How long the route has been known.

localpref—Local preference value included in the route.

from—Address and interface from which route was learned.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

AS Path—AS path included in the route.

[ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the router.

{ }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.

( )—Parentheses enclose a confederation.

([ ])—Enclose a confederation set.

Sample Output

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 568 721 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1673 1675 1747 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
inet.1: 728 destinations, 728 routes (545 active, 0 holddown, 183 hidden)
+ = Active Route, - = Last Active, * = Both
...
inet.2: 7367 destinations, 7360 routes (7307 active, 53 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
...
```

show route damping

Syntax show route damping <detail> (decayed | history | suppressed)

Description Display the routes in the routing table that have been damped.

Options none—Display all routes in the routing table that have been damped.

detail—(Optional) Display detailed information about routes in the routing table that have been damped.

decayed—Entries which are decayed but are not suppressed

history—Entries which are withdrawn but have history

suppressed—Entries suppressed due to route damping

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

Nexthop—Address of the next hop and the interface used to reach the next hop.

State—Flags for this route. It can be one or more of the following:

- Accounting—Route needs accounting.
- Active—Route is active.
- Clone—Route is a clone.
- Delete—Route has been deleted.
- Ex—Exterior route.
- Hidden—Route is not used because of routing policy.
- IfCheck—Route needs forwarding RPF check.
- Initial—Route is being added.

- Int—Interior route.
- Martian—Route is a martian.
- MartianOK—Route is exempt from martian filtering.
- NoReadvrt—Route is not to be advertised.
- NotBest—Route was not chosen because it does not have the lowest MED metric.
- NotInstall—Route is not to be installed in the forwarding table.
- Pending—Route is pending because of hold-down on another route.
- Release—Route is scheduled for release.

Local AS—AS number of the local router.

Peer AS—AS number of the peer router.

Age—How long the route has been known.

Task—Name of the task that owns the route.

AS path—AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path was originated:

- I—IGP.
- E—EGP.
- ?—Incomplete; typically, the AS path was aggregated.

Localpref—Local preference value included in the route.

Router ID—BGP router ID as advertised by the neighbor in the open message.

Merit (last update/now)—Last updated and current figure-of-merit value.

Damping parameters—Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.

Last update—Time of most recent change in path attributes.

First update—Time of first change in path attributes, which started the route damping process.

Flaps—Number of times route has gone up or down or its path attributes have changed.

Suppressed—This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.

Reusable in—Time when a suppressed route will again be available.

Preference will be—Preference value that will be applied to the route when it is again active.

Sample Output: show route damping detail

```
user@host> show route damping detail
inet.0: 21 destinations, 21 routes (15 active, 0 holddown, 6 hidden)
     + = Active Route, - = Last Active, * = Both

     1.1.1.0/24 (1 entry, 0 announced)
              BGP    Preference: /-101
                     Nexthop: 10.12.1.2 via en0.0, selected
                     State: <Hidden Ext>
                     Local AS:   911 Peer AS:   922
                     Age: 11:50
                     Task: BGP_922.10.12.1.2+179
                     AS path: 922 I
                     Localpref: 100
                     Router ID: 111.222.1.46
                     Merit (last update/now): 4833/2796
                     Default damping parameters used
                     Last update:       00:11:50
                     First update:       01:01:43
                     Flaps: 10
                     Suppressed.  Reusable in:       00:28:40
                     Preference will be: 170
...
```

Sample Output: show route damping decayed

```
user@hostname> show route damping decayed

inet.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

Sample Output: show route damping history

```
user@host> show route damping history

inet.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

Sample Output: show route damping suppressed

```
user@host> show route damping suppressed

inet.0: 30 destinations, 30 routes (30 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route exact

Syntax show route exact *destination-prefix*

Description Display only the routes that exactly match the specified address or range of addresses.

Options *destination-prefix*—Address or range of addresses.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

localpref—Local preference value included in the route.

from—Address and interface from which route was learned.

AS Path—AS path included in the route.

[ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the router.

{ }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.

( )—Parentheses enclose a confederation.

([ ])—Enclose a confederation set.

Sample Output

```
user@host> show route exact 24.226.160.0/19

inet.0: 53294 destinations, 53294 routes (53293 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

24.226.160.0/19    *[BGP/170] 00:31:04, MED 0, localpref 100, from
208.197.169.14
                   AS path: 2914 701 3493 11290 I
                    > to 111.222.5.254 via fxp0.0
```

show route export

Syntax show route export <brief | detail> <instance *instance-name* | *routing-table-name*>

Description Display a list of instances or routing tables that are participating in the export mechanism, either as importers or exporters of routes.

Options none—Display information about all routing tables.

brief—(Optional) Display brief information about instances or routing tables.

detail—(Optional) Display detailed information about instances or routing tables.

instance *instance-name*—(Optional) Display a particular routing instance where auto-export is currently enabled.

routing-table-name—Display information about a particular routing table where auto-export is currently enabled.

Required Privilege Level view

Output Fields Export—Whether the table is currently exporting routes to other tables. It displays Y or N (Yes or No).

Flags—(Detail only) Flags for this feature on this instance,

- auto-policy—The policy was deduced from the configured IGP export policies.
- cleanup—Configuration information for this instance is no longer valid.
- config—The instance was explicitly configured.

Import policy—Displays the policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf.

Instance—Name of the routing instance.

Options—Configured option displays the type of routing tables the feature handles.

- unicast—Indicates instance.inet.0.
- multicast—Indicates instance.inet.2.
- unicast multicast—Indicates instance.inet.0 and instance.inet.2.

Routes—Display the number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.

Import—Displays a list of tables currently importing routes from the originator table. Not displayed for tables that are not exporting any routes.

Table—Name of the routing tables that are either import or export routes.

Type—Type of routing instance. It can be forwarding, non-forwarding, or vrf.

Sample Output: show route export brief

```
user@host> show route export brief
Table                            Export           Routes
inet.0                           N                     0
black.inet.0                     Y                     3
red.inet.0                       Y                     4
```

Sample Output: show route export detail

```
user@host> show route export detail
inet.0                                          Routes:        0
black.inet.0                                    Routes:        3
  Import: [ inet.0 ]
red.inet.0                                      Routes:        4
  Import: [ inet.0 ]
```

Sample Output: show route export instance detail

```
user@host> show route export instance detail
Instance: master                        Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                         Type: non-forwarding
Instance: red                           Type: non-forwarding
```

show route export vrf-target

Syntax show route export vfr-target <brief | detail> <community *community-regexp*>

Description Displays the target communities for which auto-export is currently distributing routes. vrf-target is relevant when there are overlapping VPNs.

Options none—Display information about all routing tables.

brief—(Optional) Display brief information about routing tables.

detail—(Optional) Display detailed information about routing tables.

community *community-regexp*—Display information about a community with this community identifier.

Required Privilege Level view

Output Fields Route Target—Target communities for which auto-export is currently distributing routes.

Family—Routing table entries for the specified family.

Import—Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not exporting routes.

Export—Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes.

Target—(Detail only) Target communities, family, and options for which auto-export is currently distributing routes.

Options—Configured option displays the type of tables the feature handles.

- unicast—Indicates instance.inet.0.
- multicast—Indicates instance.inet.2.
- unicast multicast—Indicates instance.inet.0 and instance.inet.2.

Import table(s)—(Detail only) Name of the routing tables that are importing a particular route target.

Export table(s)—(Detail only) Name of the routing tables that are exporting a particular route target.

Sample Output: show route export vrf-target

```
user@host> show route export vrf-target
Route Target            Family                  Import     Export
69:1                    inet     unicast             2          2
69:2                    inet     unicast             2          2
```

Sample Output: show route export vrf-target community

```
user@host> show route export vrf-target community target:69:1
Route Target            Family                  Import     Export
69:1                    inet     unicast             2          2
```

Sample Output: show route export vrf-target detail

```
user@host> show route export vrf-target detail
Target: 1:12                                inet     unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13                                inet     unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0
```

show route forwarding-table

Syntax show route forwarding-table <*destination-prefix*> <family *family* | matching *matching*> <multicast> <vpn *vpn*> <detail | summary | extensive>

Description Display the route entries in the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine.

The show route forwarding-table command displays only the network-layer prefixes and their next hops. Commonly, you need information in the show route commands. However, when you run into problems, such as connectivity problems, it can be necessary to use the show route forwarding-table command to verify that the routing protocol process has relayed the correct information into the forwarding table.

> **Note** The show route forwarding-table command is an independent command, not a filter that selects specific information that is displayed from the routing tables. You cannot use this command in conjunction with any of the show route filter options (see Table 65 on page 512).

Options none—Display the routes in the forwarding table. This option is equivalent to the UNIX netstat –rn command.

destination-prefix—(Optional) IP address of a destination.

detail—(Optional) Display detailed information, including clone routes. This option is equivalent to the UNIX netstat –rna command.

family *family*—(Optional) Display routing table entries for the specified family.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

summary—(Optional) Display summary information about routes.

vpn *vpn*—(Optional) Display routing table entries for specified vpn.

Required Privilege Level view

Output Fields Routing table—Name of the routing table.

Internet, ISO, MPLS—IP, ISO, and MPLS addresses.

Destination—Destination of the route.

Route type—How the route was placed into the forwarding table (for Detail and Standard output, Type is located in the second column).

Detail, Standard	Extensive	Definition
clon	cloned	Clone route (for TCP or multicast only. Displayed only if you specify the detail option).
dest	destination	Remote addresses directly reachable through an interface.
iddn	destination down	dest route for which the interface is down.
ifcl	interface cloned	clone route for which the interface is down.
ifdn	route down	intf route for which the interface is down.
ignr	ignore	Ignore this route.
intf	interface	Installed as the result of configuring an interface.
perm	permanent	Permanent (installed by kernel when the routing table is initialized).
user	user	Installed by the routing protocol process or as a result of the configuration.

RtRef—(Detail output) Number of routes to reference.

Route reference—(Extensive output) Number of routes to reference.

InIf—(Detail output) Incoming interface.

Flags—Additional information about a route.

Standard, Detail	Extensive	Definition
0x00	none	No flags enabled.
0x01	static	Static route.
0x02	cached	Cache route.
0x04	incoming-iface *interface number*	Check against incoming interface.
0x08	accounting	Route has accounting.
0x10	sent to PFE	Route has been sent to the Packet Forwarding Engine.
0x20	prefix load balance	Load balancing has been enabled for this prefix.

Nexthop—Next hop to the destination.

Next-hop type—Information about the next hop (for Detail and Standard output, Type is located in the fifth column).

Standard, Detail	Extensive	Definition
bcst	broadcast	Broadcast.
deny	deny	Deny.
dscd	discard	Discard; no ICMP unreachable message sent.
hold	hold	Next hop is waiting to be resolved (will turn into a unicast or multicast).
idxd	indexed	Indexed next hop.
indr	indirect	Indirect next hop.
locl	local	Local address on an interface.
mcrt	routed multicast	Regular multicast next hop.
mcst	multicast	Wire multicast next hop (limited to the LAN).
mdsc	multicast discard	Multicast discard.
mgrp	multicast group	Multicast group member.
recv	receive	Receive.
rjct	reject	Discard; ICMP unreachable message sent.
rslv	resolve	Resolving next hop.
ucst	unicast	Unicast.
ulst	unilist	List of unicast next hops. A packet sent to this next hop will go to any next hop in the list.

Index—Number associated with the next hop.

NhRef—(Detail output) Number of routes that reference this next hop.

Netif—(Detail output) Interface used to reach the next hop.

Next-hop interface—(Extensive output) Interface used to reach the next hop.

RPF information—(Extensive output) List of interfaces from where the prefix can be accepted. RPF information is displayed only when rpf-check is configured on the interface.

Sample Output: show route forwarding-table (standard)

```
user@host> show route forwarding-table
Routing table: inet
Internet:
Destination        Type RtRef Nexthop           Type Index NhRef Netif
default            user     0                   dscd    10     1
default            perm     0                   rjct    12     2
2.2.2.0/24         ifdn     0 f.0.8.0           ucst    39     1
so-1/1/0.0
2.2.2.1/32         user     0                   rjct    12     2
2.2.2.1/32         intf     0 2.2.2.1           locl    38     1
4.1.1.0/24         intf     0 f.0.8.0           ucst    43     1
so-1/0/0.0
4.1.1.2/32         intf     0 4.1.1.2           locl    42     1
5.1.1.0/24         user     0                   ucst    44     4
so-1/0/0.0
10.255.245.220/32  intf     0 10.255.245.220    locl    25     1
10.255.245.245/32  user     0                   ucst    44     4
```

```
so-1/0/0.0
10.255.245.246/32  user     0                    ucst    44     4
so-1/0/0.0
192.16.0.0/12      user     1 192.168.5.254      ucst    29    11 fxp0.0
192.168.0.0/18     user     0 192.168.5.254      ucst    29    11 fxp0.0
192.168.0.0/20     user     0 192.168.5.254      ucst    29    11 fxp0.0
192.168.1.0/24     user     0 192.168.5.254      ucst    29    11 fxp0.0
192.168.5.0/24     intf     0                    rslv    22     1 fxp0.0
192.168.5.0/32     dest     0 192.168.5.0        recv    20     1 fxp0.0
192.168.5.49/32    dest     0 0:0:c0:e8:69:db    ucst    31     1 fxp0.0
192.168.5.73/32    dest     0 0:a0:c9:85:c:44    ucst    32     1 fxp0.0
192.168.5.80/32    dest     0 0:d0:b7:1e:92:f2   ucst    30     1 fxp0.0
192.168.5.220/32   intf     0 192.168.5.220      locl    21     2
192.168.5.220/32   dest     0 192.168.5.220      locl    21     2
192.168.5.254/32   dest     0 0:90:69:d:30:1     ucst    29    11 fxp0.0
192.168.5.255/32   dest     0 192.168.5.255      bcst    19     1 fxp0.0
192.168.21.0/24    user     0 192.168.5.254      ucst    29    11 fxp0.0
207.17.136.192/32  user     0 192.168.5.254      ucst    29    11 fxp0.0
207.79.80.0/24     user     0 192.168.5.254      ucst    29    11 fxp0.0
208.197.169.0/24   user     0 192.168.5.254      ucst    29    11 fxp0.0
208.223.208.0/24   user     0 192.168.5.254      ucst    29    11 fxp0.0
224.0.0.0/4        perm     1                    mdsc    11     1
224.0.0.1/32       perm     0 224.0.0.1          mcst     7     3
224.0.0.5/32       user     1 224.0.0.5          mcst     7     3
255.255.255.255/32 perm     0                    bcst     8     1

Routing table:: iso
ISO:
Destination        Type RtRef Nexthop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Nexthop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Nexthop           Type Index NhRef Netif
default            perm     0                   dscd    16     1
```

Sample Output: show route forwarding-table detail

```
user@host> show route forwarding-table detail
Internet:
Destination        Type RtRef  InIf Flags Nexthop        Type Index NhRef Netif

default            user     0     0  0x10                rjct     9     9
default            perm     0     0  0x00                rjct     9     9
1.1.1.8/32         user     0     0  0x10 111.222.8.208  ucst    35     5 so-
2/1/1.0
1.1.1.9/32         user     0     0  0x10 111.222.8.208  ucst    35     5 so-
2/1/1.0
10.0.222.0/30      ifdn     0     0  0x00 ff.3.0.21      ucst    51     1 t3-
5/2/2.0
10.0.222.1/32      user     0     0  0x10                rjct     9     9
10.0.222.1/32      intf     0     0  0x00 10.0.222.1     locl    50     1
10.1.1.76/32       user     0     0  0x10 111.222.8.208  ucst    35     5 so-
2/1/1.0
10.255.245.87/32   user     0     0  0x10 111.222.8.208  ucst    35     5 so-
2/1/1.0
```

```
12.12.12.0/24      ifdn     0     0  0x00 12.12.12.0     rjct    29     1 so-
2/1/0.1

ISO:
Destination        Type RtRef  InIf Flags Nexthop        Type Index NhRef Netif
default            perm     0     0  0x00                dscd    17     1
47.0005.80ff.f800.0000.0108.0001.1921.6800.4021.00
                   intf     0     0  0x00                locl    18     1

MPLS:
Interface.Label    Type RtRef  InIf Flags Nexthop        Type Index NhRef Netif
default            perm     0     0  0x00                dscd     1     1
0000000            user     0     0  0x18                recv     3     2
0000001            user     0     0  0x18                recv     3     2
```

Sample Output: show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination:  default
  Type: user                RtRef: 0
  Flags: sent to PFE
  Next-hop type: discard               index: 4        refcnt: 2

Destination:  default
  Type: permanent           RtRef: 0
  Flags: none
  Next-hop type: reject                index: 6        refcnt: 3

Destination:  10.10.12.0/24
  Type: interface           RtRef: 0
  Flags: sent to PFE
  Next-hop type: resolve               index: 37       refcnt: 1
  Netif: fe-2/3/0.0

Destination:  10.10.12.0/32
  Type: destination         RtRef: 0
  Flags: sent to PFE
  Next-hop: 10.10.12.0
  Next-hop type: receive               index: 35       refcnt: 1
  Netif: fe-2/3/0.0
  RPF Information: fe-2/3/0.0 so-3/0/0.0

Destination:  224.0.0.0/4
  Type: permanent           RtRef: 1
  Flags: sent to PFE
  Next-hop type: multicast discard     index: 5        refcnt: 2

Destination:  224.0.0.1/32
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop: 224.0.0.1
  Next-hop type: multicast             index: 1        refcnt: 4

Destination:  224.0.0.5/32
  Type: user                RtRef: 1
  Flags: sent to PFE
  Next-hop: 224.0.0.5
  Next-hop type: multicast             index: 1        refcnt: 4
```

```
Destination:  255.255.255.255/32
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop type: broadcast             index: 2        refcnt: 2

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop type: reject                index: 14       refcnt: 2

Destination:  fe00::/8
  Type: route down          RtRef: 0
  Flags: none
  Next-hop type: resolve               index: 40       refcnt: 1
  Netif: fe-2/3/1.0

Destination:  fe08::1/128
  Type: user                RtRef: 0
  Flags: sent to PFE
  Next-hop type: reject                index: 14       refcnt: 2

Destination:  fe08::1/128
  Type: interface           RtRef: 0
  Flags: none
  Next-hop: fe08::1
  Next-hop type: local                 index: 38       refcnt: 2

Destination:  fe08::1/128
  Type: destination down    RtRef: 0
  Flags: none
  Next-hop: fe08::1
  Next-hop type: local                 index: 38       refcnt: 2

Destination:  ff00::/8
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop type: multicast discard     index: 13       refcnt: 1

Destination:  ff02::1/128
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop: ff02::1
  Next-hop type: multicast             index: 10       refcnt: 1

Destination:  ff02::1:ff00:1/128
  Type: route down          RtRef: 0
  Flags: none
  Next-hop: ff02::1:ff00:1
  Next-hop type: receive               index: 39       refcnt: 1
  Netif: fe-2/3/1.0

Routing table:: ccc [Index 0]
MPLS:

Destination:  default
  Type: permanent           RtRef: 0
  Flags: sent to PFE
  Next-hop type: discard               index: 17       refcnt: 1
```

show route inactive

Syntax show route inactive

Description Display the inactive entries in each routing table.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

via—Interface to reach the next hop.

Sample Output

```
user@host> show route inactive
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 19:40:17
                    > via lo0.0
```

show route instance

Syntax show route instance <detail | summary><*instance-name*>

Description Display routing instance information.

Options detail—(Optional) Display detailed routing instance information.

instance-name—(Optional) Display routing instance information for specified instance.

summary—(Optional) Display summary routing instance information.

Required Privilege Level view

Output Fields Instance—Name of routing instance.

Type—Type of routing instance. It can be forwarding, no-forwarding, or vrf.

State—State of the routing instance. It can be active or inactive.

Interfaces—Name of interfaces belonging to this routing instance.

Route-distinguisher—The unique route distinguisher associated with this routing instance.

Vrf-import—VPN routing and forwarding instance import policy name.

Vrf-export—VPN routing and forwarding instance export policy name.

Tables—Tables associated with this routing instance.

Primary rib—Primary table for this routing instance.

Active/holddown/hidden—Number of active, holddown, and hidden routes.

Sample Output: show route instance summary

```
user@host> show route instance summary
Instance           Type                Primary rib      Active/holddown/hidden
yellow-vpn         vrf                 yellow-vp.inet.0 5/0/0
pink-vpn           vrf                 pink-vpn.inet.0  4/0/0
green-vpn          vrf                 green-vpn.inet.0 6/0/0
blue-vpn           vrf                 blue-vpn.inet.0  5/0/0
master             forwarding          inet.0           29/0/1
```

Sample Output: show route instance detail

```
user@host> show route instance detail
yellow-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/3.0
  Route-distinguisher: 69:40
  Vrf-import: [ yellow-vpn-import ]
  Vrf-export: [ yellow-vpn-export ]
  Tables:
    yellow-vp.inet.0       : 5 routes (5 active, 0 holddown, 0 hidden)
pink-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/1.0
  Route-distinguisher: 69:30
  Vrf-import: [ pink-vpn-import ]
  Vrf-export: [ pink-vpn-export ]
  Tables:
    pink-vpn.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
green-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/2.0
  Route-distinguisher: 69:20
  Vrf-import: [ green-vpn-import ]
  Vrf-export: [ green-vpn-export ]
  Tables:
    green-vpn.inet.0       : 6 routes (6 active, 0 holddown, 0 hidden)
blue-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/0.0
  Route-distinguisher: 69:10
  Vrf-import: [ blue-vpn-import ]
  Vrf-export: [ blue-vpn-export ]
  Tables:
    blue-vpn.inet.0        : 5 routes (5 active, 0 holddown, 0 hidden)
master:
  Type: forwarding        State: Active
  Tables:
    inet.0                 : 30 routes (29 active, 0 holddown, 1 hidden)
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    mpls.0                 : 9 routes (9 active, 0 holddown, 0 hidden)
```

show route label-switched-path

Syntax show route label-switched-path *path-name*

Description Display the routes used in a label-switched path.

Options *path-name*—Name of a label-switched path.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

metric—Metric associated with the route.

to—Destination of the LSP; that is, the address of the egress router.

via—Interface to reach the LSP.

lsp-name—Name of the LSP.

Sample Output

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                      to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route martians

Syntax show route martians <table *routing-table*>

Description Display the martian entries associated with each routing table.

> **Note** The show route martians command is an independent command, not a filter that selects specific information that is displayed from the routing tables. You cannot use this command in conjunction with any of the show route filter options (see Table 65 on page 512).

Options table *routing-table*—(Optional) Display the martian entries associated with a particular routing table.

Required Privilege Level view

Sample Output

```
user@host> show route martians
inet.0:
0.0.0.0/0          Match: Exact, Advertise: <none>
0.0.0.0/8          Match: <none>, Advertise: <none>
127.0.0.0/8        Match: <none>, Advertise: Restrict
128.0.0.0/16       Match: <none>, Advertise: Restrict
191.255.0.0/16     Match: <none>, Advertise: Restrict
192.0.0.0/24       Match: <none>, Advertise: Restrict
223.255.255.0/24   Match: <none>, Advertise: Restrict
240.0.0.0/4        Match: <none>, Advertise: Restrict

inet.1:
0.0.0.0/0          Match: Exact, Advertise: <none>
0.0.0.0/8          Match: <none>, Advertise: <none>
127.0.0.0/8        Match: <none>, Advertise: Restrict
128.0.0.0/16       Match: <none>, Advertise: Restrict
191.255.0.0/16     Match: <none>, Advertise: Restrict
192.0.0.0/24       Match: <none>, Advertise: Restrict
223.255.255.0/24   Match: <none>, Advertise: Restrict
240.0.0.0/4        Match: <none>, Advertise: Restrict

inet.2:
0.0.0.0/0          Match: Exact, Advertise: <none>
0.0.0.0/8          Match: <none>, Advertise: <none>
127.0.0.0/8        Match: <none>, Advertise: Restrict
128.0.0.0/16       Match: <none>, Advertise: Restrict
191.255.0.0/16     Match: <none>, Advertise: Restrict
192.0.0.0/24       Match: <none>, Advertise: Restrict
223.255.255.0/24   Match: <none>, Advertise: Restrict
240.0.0.0/4        Match: <none>, Advertise: Restrict

inet.3:
0.0.0.0/0          Match: Exact, Advertise: <none>
0.0.0.0/8          Match: <none>, Advertise: <none>
127.0.0.0/8        Match: <none>, Advertise: Restrict
128.0.0.0/16       Match: <none>, Advertise: Restrict
191.255.0.0/16     Match: <none>, Advertise: Restrict
192.0.0.0/24       Match: <none>, Advertise: Restrict
223.255.255.0/24   Match: <none>, Advertise: Restrict
240.0.0.0/4        Match: <none>, Advertise: Restrict
```

show route next-hop

Syntax show route next-hop *destination-prefix*

Description Display the entries in the routing table that are being sent to the specified next-hop address.

Options *destination-prefix*—IP address or name of a next hop.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

to—Next hop to the destination.

via—Interface to reach the next hop.

Sample Output

```
user@host> show route next-hop 111.222.5.254

inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:43:39
                    > to 111.222.5.254 via fxp0.0
```

show route output

Syntax show route output (interface *interface* | address *address*)

Description Display the entries in the routing table that are to be sent out the specified interface or out the interface with the specified address.

Options address *address*—Interface address.

interface *interface*—Interface name.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

via—Interface to reach the next hop.

Sample Output

```
user@host> show route output interface t3-5/2/1

inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.13.13.13/32     *[Direct/0] 00:44:55
                    > via t3-5/2/1.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route protocol

Syntax show route protocol *protocol*

Description Display the route entries in the routing table that were learned from a particular protocol.

Options *protocol*—Protocol from which the route was learned. It can be one of the following: aggregate, bgp, direct, dvmrp, isis, ldp, local, mpls, msdp, ospf, pim, rip, ripng, rsvp, or static.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

via—Interface to reach the next hop.

Sample Output

```
user@host> show route protocol direct
inet.0: 35 destinations, 35 routes (34 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 14:36:24
                     > via lo0.0
111.222.5.0/24     *[Direct/0] 14:36:24
                     > via fxp0.0
111.222.8.16/28    *[Direct/0] 14:36:24
                     > via at-5/3/0.0
111.222.8.100/30   *[Direct/0] 14:36:24
                     > via at-5/3/0.129
111.222.8.104/30   *[Direct/0] 14:36:24
                     > via at-5/3/0.128
111.222.8.161/32   *[Direct/0] 14:36:24
                     > via t3-5/2/0.0
111.222.8.163/32   *[Direct/0] 14:36:24
                     > via t3-5/2/1.0
...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.5081.00/160
                   *[Direct/0] 14:36:24
                     > via lo0.0
```

show route range

Syntax show route range *destination-prefix*

Description Display all the routing table entries in a prefix range.

Options *destination-prefix*—Destination prefix and mask for the range.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

metric—Metric associated with the route.

tag—Tag associated with the route.

to—Next hop to the destination.

via—Interface to reach the next hop.

Sample Output

```
user@host> show route range
inet.0: 27 destinations, 27 routes (26 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 06:03:18
                     Discard
10.255.245.87/32   *[Direct/0] 06:03:19
                    > via lo0.0
123.16.1.0/24      *[Static/5] 06:03:18
                    > to 111.222.4.254 via fxp0.0
123.16.14.0/24     *[Static/5] 06:03:18
                    > to 111.222.4.254 via fxp0.0
```

```
                    123.16.100.0/24    *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.16.128.0/19    *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.16.160.0/24    *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.17.2.0/24      *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.17.12.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.17.16.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.17.17.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    123.17.20.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    111.222.1.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    111.222.4.0/24     *[Direct/0] 06:03:19
                                        > via fxp0.0
                    111.222.4.18/32    *[Local/0] 06:03:19
                                         Local
                    111.222.5.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    111.222.8.196/32   *[Local/0] 06:02:12
                                         Local
                    111.222.8.197/32   *[Direct/0] 06:01:54
                                        > via so-2/1/3.0
                    111.222.8.206/32   *[Direct/0] 06:01:55
                                        > via so-2/1/0.0
                    111.222.8.207/32   *[Local/0] 06:02:12
                                         Local
                    111.222.8.210/32   *[Direct/0] 05:57:06
                                        > via so-2/1/2.0
                    111.222.8.211/32   *[Local/0] 06:02:12
                                         Local
                    111.222.8.217/32   *[Local/0] 06:02:12
                                         Reject
                    124.79.80.0/24     *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    125.197.169.0/24   *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
                    125.223.208.0/24   *[Static/5] 06:03:18
                                        > to 111.222.4.254 via fxp0.0
```

Sample Output

```
user@host> show route range 123.16/16
inet.0: 27 destinations, 27 routes (26 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

123.16.1.0/24      *[Static/5] 06:03:27
                    > to 111.222.4.254 via fxp0.0
123.16.14.0/24     *[Static/5] 06:03:27
                    > to 111.222.4.254 via fxp0.0
123.16.100.0/24    *[Static/5] 06:03:27
                    > to 111.222.4.254 via fxp0.0
123.16.128.0/19    *[Static/5] 06:03:27
                    > to 111.222.4.254 via fxp0.0
123.16.160.0/24    *[Static/5] 06:03:27
                    > to 111.222.4.254 via fxp0.0
```

show route receive-protocol

Syntax show route receive-protocol *protocol neighbor-address* <detail | extensive | standard>

Description Display the routing information as it was received through a particular neighbor of a particular dynamic routing protocol. This information includes the routes that the local router advertised to the neighbor. The information displayed reflects the routes before they are filtered by that protocol's import routing policy statements and is placed into the routing table.

Options none—Display information that was received from the neighbor.

detail—(Optional) Display additional information that was received from the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were received, including communities, route reflection, and aggregator.

extensive—(Optional) Display extensive information that was received from the neighbor.

neighbor-address—Address of the neighboring router from which the route entry was received.

protocol—Protocol transmitting the route. It can be one of the following: bgp, dvmrp, msdp, pim, rip or ripng.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table, for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Next hop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

Route Distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.

VPN Label—Label assigned to this VPN.

Communities—Community path attribute of the route.

AS path: I <Originator>—(For route reflected output only) Route reflector set the originator ID attribute.

Cluster list—(For router reflected output only) Cluster id sent by the route reflector.

Originator ID—(For router reflected output only) Address of router that originally sent route to the route reflector.

Sample Output

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive

     inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop                MED    Lclpref AS path
     1.1.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     165.3.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     165.4.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     195.1.2.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68

     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route resolution

Syntax show route resolution <*prefix*> <table *table-name*> <unresolved> <brief | detail | summary>

Description Display the entries in the next-hop resolution database. The next-hop resolution database provides for recursive resolution of next hops through other prefixes in the Routing Table.

Options brief—(Default) Display brief information about entries in the resolution database.

detail—(Optional) Display detailed information about entries in the resolution database.

prefix—(Optional) Display database entries for the specified address. Specify the address as *network/destination-prefix*.

summary—(Optional) Display summary information about entries in the resolution database.

table *table-name*—(Optional) Display information about a particular table.

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level view

Output Fields Table—Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.

Originating RIB—Name of the Routing Table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving via inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.

Metric—Metric associated with the forwarding next hop.

Forwarding nexthops—Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Next hop type—(Detail output only) Type of the next hop. It can be Discard, Local, Interface, or Reject.

Indirect nexthops—An index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.

Nexthop—Number of forwarding next-hops.The Forwarding Next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Has inactive paths with indirect nexthops—(Detail output only) When a prefix is learned from multiple paths, this field precedes the list of inactive paths that have route resolution data associated with them.

protocol nexthop—Network address of a remote router that advertised the prefix. This address is used to recurse through and derive a forwarding next hop.

Indirect path forwarding nexthops—Number of indirect-path-forwarding next hops, followed by a list of associated forwarding next hops.

Sample Output: show route resolution

```
user@host> show route resolution detail
Table inet.0 Nodes 111590
0.0.0.0/0 Originating RIB: inet.0
  Metric: 0  Forwarding nexthops: 0
        Next hop type: Discard
1.0.0.0/8 Originating RIB: inet.0
  Metric: 2  Indirect nexthops: 1
        Protocol nexthop: 10.255.14.181 Metric: 2 Indirect nexthop: 84cc770 104
        Indirect path forwarding nexthops: 1
                Nexthop: 10.19.2.2 via t1-0/2/0.0
        10.255.14.181/32 Originating RIB: inet.0
          Metric: 2       Forwarding nexthops: 1
                Nexthop: 10.19.2.2 via t1-0/2/0.0
  Has inactive paths with indirect nexthops:
  Indirect nexthops: 1
        Protocol nexthop: 10.255.14.175 Metric: 4 Indirect nexthop: 84cc6e8 103
        Indirect path forwarding nexthops: 1
                Nexthop: 10.19.4.2 via t3-0/3/0.0
        10.255.14.175/32 Originating RIB: inet.0
          Metric: 4       Forwarding nexthops: 1
                Nexthop: 10.19.4.2 via t3-0/3/0.0
...
```

show route source-gateway

Syntax show route source-gateway *address*

Description Display the entries in the routing table that were learned from a particular address. The Source field in the show route detail command output lists the source for each route, if known.

Options *address*—IP address of the system.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

from—Address and interface from which route was learned.

AS Path—AS path included in the route.

[ ]—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the router.

{ }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.

( )—Parentheses enclose a confederation.

([ ])—Enclose a confederation set.

Sample Output

```
user@host> show route source-gateway 131.103.20.49
inet.0: 46035 destinations, 46035 routes (46033 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 04:43:40, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                     to 192.156.169.1 via 192.156.169.14(de0.0)
6.0.0.0/8          *[BGP/170] 04:43:40, from 131.103.20.49
                   AS Path: {666} 234 2548 568 721 Incomplete
                     to 192.156.169.1 via 192.156.169.14(de0.0)
...
```

show route summary

Syntax show route summary

Description Display summary statistics about the entries in the routing table.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Direct—Routes on the directly connected network.

Local—Local routes.

protocol-name—Name of protocol from which the route was learned.

Sample Output

```
user@host> show route summary
inet.0: 66452 destinations, 66452 routes (66452 active, 0 holddown, 0 hidden)
              Direct:      3 routes,      3 active
               Local:      2 routes,      2 active
                 BGP:  84654 routes,  64536 active
               IS-IS:   1904 routes,   1903 active

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Static:      1 routes,      1 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
              Direct:      1 routes,      1 active
```

show route table

Syntax show route table *routing-table-name* <detail | extensive>

Description Display the route entries in a particular routing table.

Options none—Display detailed information about route entries in all routing tables.

detail—(Default) Display detailed information about route entries in a particular routing table.

extensive—(Default) Display very detailed information about route entries in a particular routing table.

routing-table-name—Name of a routing table, in the format inet.*number*.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

label stacking—On the penultimate-hop router, indicates the depth of the MPLS label stack, where the label popping operation is needed. A pair of routes is displayed because the pop operation is different depending on the stack depth.

- S=0 route indicates the packet with an incoming label stack depth n >=2 exits this router using MPLS, with *n*–1 stack label.
- No S= information indicates a normal MPLS route, which has a stack depth of 1.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

Nexthop—(Detail and Extensive output only) Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Action—(Detail and Extensive output only) Can be pop, where a label is popped, or swap, where a label is replaced by another label.

State—(Detail and Extensive output only) Flags for this route. It can be one or more of the following:

- Accounting—Route needs accounting.
- Active—Route is active.
- Clone—Route is a clone.
- Delete—Route deleted.
- Ex—Exterior route.
- Hidden—Route not used because of routing policy.
- IfCheck—Route needs forwarding RPF check.
- Initial—Route being added.
- Int—Interior route.
- Martian—Route is a martian.
- MartianOK—Route exempt from martian filtering.
- NoReadvrt—Route not to be advertised.
- NotBest—Route not chosen because it does not have the lowest MED metric.
- NotInstall—Route not to be installed in the forwarding table.
- Pending—Route pending because of holddown on another route.
- Release—Route scheduled for release.
- Secondary—Route not a primary route.

Metric 1—First metric value in the route.

Task—Name of the protocol that has added the route.

Announcement bits—(Detail and Extensive output only) List of protocols that announce this route. n-Resolve inet.x indicates that the route is used for route resolution for next hops found in the routing table itnet.x. n is an index used by Juniper Networks Customer Support only.

via—Interface to reach the next hop.

Prefixes bound to route—(Detail and Extensive output only) The Forwarding Equivalent Class (FEC) that is bound to this route. This field is applicable to routes installed by LDP only.

Sample Output: show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                      Local
12.12.12.21/32     *[Local/0] 00:51:57
                      Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                      Local
13.13.13.21/32     *[Local/0] 00:51:58
                      Local
13.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                      Local
224.0.0.5/32       *[OSPF/10] 00:51:58, metric 1
```

Sample Output: show route table mpls

```
user@host> show route table mpls
mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0           *[MPLS/0] 6d 05:37:09, metric 1
                      Receive
1           *[MPLS/0] 6d 05:37:09, metric 1
                      Receive
100061     *[RSVP/7] 00:20:59, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path dc.1
100061(S=0)   *[RSVP/7] 00:20:59, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path dc.1
100063     *[RSVP/7] 00:20:58, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path to-sj2-3
100063(S=0)   *[RSVP/7] 00:20:58, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path to-sj2-3
100073     *[RSVP/7] 00:04:45, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path to-sj2-1
100073(S=0)   *[RSVP/7] 00:04:45, metric 1
           > to 192.168.200.26 via fxp1.0, label-switched-path to-sj2-1
```

Sample Output: show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
                >via fe-0/1/0.0
fec0:0:0:3::/128 *[Local/0] 00:01:34
                >Local
fec0:0:0:4::/64 *[Static/5] 00:01:34
                >to fec0:0:0:3::ffff via fe-0/1/0.0
```

Sample Output: show route table inet6.3

```
user@router> show route table inet6.3

inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

Sample Output: show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}

        *LDP    Preference:   9
                Nexthop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32


100001 (1 entry, 1 announced)
TSI:
KRT in-kernel 100001 /36 -> {so-1/0/0.0}

        *LDP    Preference:   9
                Nexthop:  via so-1/0/0.0, selected
                Swap 100007
                State: <Active Int>
                Age: 27:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 1.1.1.1/32
                                         2.2.2.2/32
```

show route terse

Syntax show route terse

Description Display a high-level summary of the routes in the routing table.

Required Privilege Level view

Output Fields A—Active route. An asterisk (*) indicates that this is the active route.

Destination—Destination of the route.

P—Protocol through which the route was learned:

- S—Static
- D—Direct
- L—Local
- I—IS-IS
- O—OSPF

Prf—Preference value of the route.

Metric 1—First metric value in the route.

Metric 2—Second metric value in the route.

Next hop—Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.

AS path—AS path included in the route.

Sample Output

```
user@host> show route terse
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf Metric 1   Metric 2    Next hop        AS path
* 0.0.0.0/0          S   5                       >111.222.5.254
* 1.0.0.1/32         D   0                       >at-5/3/0.0
* 1.0.0.2/32         L   0                        Local
* 12.12.12.21/32     L   0                        Reject
* 13.13.13.13/32     D   0                       >t3-5/2/1.0
* 13.13.13.14/32     L   0                        Local
* 13.13.13.21/32     L   0                        Local
* 13.13.13.22/32     D   0                       >t3-5/2/0.0
  127.0.0.1/32       D   0                       >lo0.0
* 111.222.5.0/24     D   0                       >fxp0.0
* 111.222.5.81/32    L   0                        Local
* 224.0.0.5/32       O  10          1             MultiRecv
```

Chapter 35
IS-IS Monitoring and Troubleshooting

Table 67 summarizes the command-line interface (CLI) commands you can use to monitor IS-IS. In the table, the commands are grouped by functionality. In the remainder of this chapter, the commands are explained alphabetically.

Table 67: Commands for Monitoring IS-IS

Task or Information to Monitor	Command
Status of interfaces on which IS-IS is running	show isis interface on page 584
Hostname mapping	show isis hostname on page 583
IS-IS routing table entries	show isis routes on page 587
Database entries	show isis database on page 580
Remove database entries	clear isis database on page 575
Adjacent routers	show isis adjacency on page 577
Remove adjacencies	clear isis adjacency on page 574
SPF calculations	show isis spf on page 588
IS-IS traffic statistics	show isis statistics on page 591
Zero IS-IS traffic statistics	clear isis statistics on page 576

System ID Format in Command Output

In IS-IS command output, the CLI displays the system identifier (sysid) numerically by default. If you want the host name to be displayed instead, you must define the sysid and map it to a host name. To do this, include the static-host-mapping statement at the [edit system] hierarchy level of the configuration:

```
[edit system]
cli# show
static-host-mapping{
    host-name {
        sysid system-identifier;
    }
}
```

clear isis adjacency

Syntax clear isis adjacency <*system-id*>

Description Remove the entries from the IS-IS adjacency database.

Options none—Remove all entries from the adjacency database.

system-id—(Optional) Remove the specified entry from the adjacency database. Specify *system-id* by name, not by address prefix.

Required Privilege Level clear

See Also show isis adjacency on page 577

Sample Output

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Up                   26
so-1/1/3.0    1921.6800.5080 3 Up                   23
so-5/0/0.0    1921.6800.5080 3 Up                   19
user@host> clear isis adjacency karakul
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing         26
so-1/1/3.0    1921.6800.5080 3 Up                   24
so-5/0/0.0    1921.6800.5080 3 Up                   21
```

clear isis database

Syntax clear isis database <*system-id*>

Description Remove the entries from the IS-IS link-state database, which contains data about protocol data unit (PDU) packets.

Options none—Remove all the entries from the link-state database.

system-id—(Optional) Remove the specified entries. Specify *system-id* by name, not by address prefix.

Required Privilege Level clear

See Also show isis database on page 580

Sample Output

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x12   0x84dd            1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
crater.00-00                0x19   0xe92c            1134
badlands.00-00              0x16   0x1454             985
carlsbad.00-00              0x33   0x220b            1015
ranier.00-00                0x2e   0xfc31            1007
1921.6800.5066.00-00        0x11   0x7313             566
1921.6800.5067.00-00        0x14   0xd9d4             939
  6 LSPs
user@host> clear isis database
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime (secs)
```

clear isis statistics

Syntax clear isis statistics

Description Zero statistics about IS-IS traffic.

Required Privilege Level view

See Also show isis statistics on page 591

Sample Output

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12793      12793          0       8666        719
IIH            116751     116751          0     118834          0
CSNP           203956     203956          0     204080          0
PSNP             7356       7350          6       8635          0
Unknown             0          0          0          0          0
Totals         340856     340850          6     340215        719

Total packets received: 340856 Sent: 340934

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1064
Fragments rebuilt:       1087
LSP regenerations:        436
Purges initiated:           0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP                 0          0          0          0          0
IIH                 3          3          0          3          0
CSNP                2          2          0          4          0
PSNP                0          0          0          0          0
Unknown             0          0          0          0          0
Totals              5          5          0          7          0

Total packets received: 5 Sent: 7

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                   0
Fragments rebuilt:          0
LSP regenerations:          0
Purges initiated:           0
```

show isis adjacency

Syntax show isis adjacency <*system-id*> <brief | detail | extensive>

Description Display information about IS-IS neighbors.

Options brief—(Optional) Display brief information about the entries in the IS-IS link-state database.

detail—(Optional) Display detailed information about the entries in the IS-IS link-state database. Entries are listed by system identifier (sysid).

extensive—(Optional) Display extensive information about the entries in the IS-IS link-state database, including the time and reason for recent transitions.

system-id—(Optional) Display the entries in the link-state database for the specified system.

Default: brief

Required Privilege Level view

See Also clear isis adjacency on page 574

Sample Output Sample Output: show isis adjacency brief on page 578
Sample Output: show isis adjacency detail on page 579
Sample Output: show isis adjacency extensive on page 579

Options at a Glance Table 68 summarizes which information is included in each show isis adjacency command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 68: Show IS-IS Adjacency Output Field Summary

Options	Field Description
Detail Extensive	Circuit type—Bit mask of levels on this interface. L1 = Level 1 router, L2 = Level 2 router, L1/L2 = both Level 1 and Level 2 router.
Detail Extensive	Expires in—How long until the adjacency expires, in seconds.
Brief	Hold (secs)—Remaining hold time of the adjacency.
All	Interface—Interface through which the neighbor is reachable.
Detail Extensive	IP addresses—IP address of this neighbor.
Detail Extensive	Last transition—Time of the last Up/Down transition.
All	L or Level—Level.
Detail Extensive	Restart capable—Neighbor is capable of graceful restart.
Detail Extensive	Priority—Priority to become the designated intermediate system.
Brief	SNPA—Subnetwork point of attachment (MAC address of the next hop).
Detail Extensive	Speaks—Protocols supported by this neighbor.
All	State—State of the adjacency. It can be Up, Down, New, Initializing, or Rejected.
All	System—System identifier (sysid), printed as a name if possible.
Extensive	Transition log—List of recent transitions, including When, State, and Reason.
Detail Extensive	Up/Down transitions—Count of adjacency status changes from Up to Down or from Down to Up.

Output Fields

Interface—Interface through which the neighbor is reachable.

System—(Brief output only) System identifier (sysid), printed as a name if possible.

L or Level—Level:

- 1—Level 1 only
- 2—Level 2 only
- 3—Level 1 and Level 2

An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.

State—State of the adjacency. It can be Up, Down, New, One-way, Initializing, or Rejected.

Hold (secs)—(Brief output only) Remaining hold time of the adjacency.

SNPA—(Brief output only) Subnetwork point of attachment (MAC address of the next hop).

Expires in—(Detail output only) How long until the adjacency expires, in seconds.

Priority—(Detail output only) Priority to become the designated intermediate system.

Up/Down transitions—(Detail output only) Count of adjacency status changes from Up to Down or from Down to Up.

Last transition—(Detail output only) Time of the last Up/Down transition.

Circuit type—(Detail output only) Bit mask of levels on this interface. L1 = Level 1 router, L2 = Level 2 router, L1/L2 = both Level 1 and Level 2 router.

Restart capable—Neighbor is capable of graceful restart. The value can be Yes or No.

Speaks—(Detail output only) Protocols supported by this neighbor.

IP addresses—(Detail output only) IP address of this neighbor.

Transition log—(Extensive output only) List of recent transitions, including When, State, and Reason.

Sample Output: show isis adjacency brief

```
user@host> show isis adjacency brief
IS-IS adjacency database:
Interface  System         L State        Hold (secs) SNPA
so-1/0/0.0 1921.6800.5067 2 Up                   13
so-1/1/0.0 1921.6800.5067 2 Up                   25
so-1/2/0.0 1921.6800.5067 2 Up                   20
so-1/3/0.0 1921.6800.5067 2 Up                   19
so-2/0/0.0 1921.6800.5066 2 Up                   19
so-2/1/0.0 1921.6800.5066 2 Up                   17
so-2/2/0.0 1921.6800.5066 2 Up                   20
so-2/3/0.0 1921.6800.5066 2 Up                   20
so-5/0/0.0 ranier         2 Up                   17
```

Sample Output: show isis adjacency detail

```
user@host> show isis adjacency detail
IS-IS adjacency database:

1921.6800.5067
  Interface: so-1/0/0.0, Level: 2, State: Up, Expires in: 17 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 02:18:14 ago
  Circuit type: 0, Speaks: IP
  Restart capable: Yes
  IP addresses: 192.168.129.34

1921.6800.5067
  Interface: so-1/1/0.0, Level: 2, State: Up, Expires in: 20 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 02:18:13 ago
  Circuit type: 0, Speaks: IP
  Restart capable: Yes
  IP addresses: 192.168.129.38

1921.6800.5067
  Interface: so-1/2/0.0, Level: 2, State: Up, Expires in: 24 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 02:18:23 ago
  Circuit type: 0, Speaks: IP
  Restart capable: Yes
  IP addresses: 192.168.129.42
...
```

Sample Output: show isis adjacency extensive

```
user@host> show isis adjacency extensive
  isis2
    Interface: sr1.0, Level: 3, State: Up, Expires in 25 secs
    Priority: 0, Up/Down transitions: 16, Last transition: 00:00:01 ago
    Circuit type: 3, Speaks: IP
    Restart capability: Yes
    IP addresses: 123.456.36.18
    Transition log:
    When                  State        Reason
    Fri Jun 29 17:32:26   Up           Seenself
    Fri Jun 29 17:32:42   Down         Interface Down
    Fri Jun 29 17:32:52   Up           Seenself
    Fri Jun 29 17:33:19   Down         Level Mismatch
    Fri Jun 29 17:33:25   Up           Seenself
    Fri Jun 29 17:33:43   Down         Level Mismatch
    Fri Jun 29 17:33:49   Up           Seenself
    Fri Jun 29 17:35:40   Rejected     Error
    Fri Jun 29 17:36:12   Up           Seenself
    Fri Jun 29 17:36:47   Down         Interface Down
    Fri Jun 29 17:36:59   Up           Seenself
    Fri Jun 29 17:37:07   Down         Interface Down
    Fri Jun 29 17:37:19   Up           Seenself
    Fri Jun 29 17:37:27   Down         Interface Down
    Fri Jun 29 17:37:38   Up           Seenself
    Fri Jun 29 17:37:55   Down         Interface Down
```

show isis database

Syntax show isis database <*system-id*> <brief | detail | extensive>

Description Display the entries in the IS-IS link-state database, which contains data about PDU packets.

Options brief—(Optional) Display brief information about the entries in the link-state database.

detail—(Optional) Display detailed information about the entries in the link-state database.

extensive—(Optional) Display very detailed information about the entries in the link-state database.

system-id—(Optional) Display the link-state database entries for the specified system.

Default: brief

Required Privilege Level view

See Also clear isis database on page 575

Sample Output Sample Output: show isis database brief on page 581
Sample Output: show isis database detail on page 581
Sample Output: show isis database extensive on page 582

Options at a Glance Table 69 summarizes which information is included in each show isis database command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 69: Show ISIS Database Output Field Summary

Options	Field Description
Extensive	Additional fields contain internal IS-IS information that is useful only in troubleshooting obscure problems.
All	Checksum—Checksum value of the LSP.
Detail Extensive	IP prefix—Prefix advertised by this LSP.
Detail	IS neighbor—An IS-IS neighbor of the advertising system.
All	Lifetime (secs)—Remaining lifetime of the LSP, in seconds.
All	LSP ID—Link-state PDU (LSP) identifier.
Detail Extensive	LSP-id—Link-state PDU (LSP) identifier.
Detail Extensive	Metric—Metric of the prefix or neighbor.
All	Sequence—Sequence number of the LSP.

Output Fields

LSP ID—Link-state PDU (LSP) identifier.

Sequence—Sequence number of the LSP.

Checksum—Checksum value of the LSP.

Lifetime (secs)—Remaining lifetime of the LSP, in seconds.

IP prefix—(Detail and Extensive output only) Prefix advertised by this LSP.

IS neighbor—(Detail output only) An IS-IS neighbor of the advertising system.

Metric—(Detail and Extensive output only) Metric of the prefix or neighbor.

The remaining fields contain internal IS-IS information that is useful only in troubleshooting obscure problems. For more details about these fields, contact your customer support representative.

Sample Output: show isis database brief

```
user@host> show isis database brief
IS-IS level 1 link-state database:
LSP ID                      Sequence Checksum Lifetime (secs)
1921.6800.5080.00-00            0x89    0xfad             682
1921.6800.5081.00-00            0x69   0x8b2b            1103
  2 LSPs

IS-IS level 2 link-state database:
LSP ID                      Sequence Checksum Lifetime (secs)
1921.6800.5080.00-00            0x8d    0x7b1             685
1921.6800.5081.00-00            0x69   0x8b2b            1063
  2 LSPs
```

Sample Output: show isis database detail

```
user@host> show isis database detail
IS-IS level 1 link-state database:
1921.6800.5080.00-00  Sequence: 0x89, Checksum: 0xfad, Lifetime: 1198 secs
  IS neighbor:    1921.6800.5080.01 Metric:     10
  IS neighbor:    1921.6800.5081.00 Metric:     10
  IS neighbor:    1921.6800.5080.02 Metric:      0
  IP prefix:       192.168.8.198/32 Metric:     10 Internal
  IP prefix:       192.168.8.192/32 Metric:     10 Internal
  IP prefix:       192.168.8.196/32 Metric:     10 Internal
1921.6800.5081.00-00  Sequence: 0x69, Checksum: 0x8b2b, Lifetime: 1199 secs
  IS neighbor:    1921.6800.5080.00 Metric:     10
  IS neighbor:    1921.6800.5081.02 Metric:      0
  IP prefix:       192.168.8.197/32 Metric:     10 Internal
  IP prefix:       192.168.8.193/32 Metric:     10 Internal

IS-IS level 2 link-state database:
1921.6800.5080.00-00  Sequence: 0x8d, Checksum: 0x7b1, Lifetime: 1198 secs
  IS neighbor:    1921.6800.5080.01 Metric:     10
  IS neighbor:    1921.6800.5081.00 Metric:     10
  IS neighbor:    1921.6800.5080.02 Metric:      0
  IP prefix:       192.168.8.198/32 Metric:     10 Internal
  IP prefix:       192.168.8.192/32 Metric:     10 Internal
  IP prefix:       192.168.8.196/32 Metric:     10 Internal
1921.6800.5081.00-00  Sequence: 0x69, Checksum: 0x8b2b, Lifetime: 1199 secs
  IS neighbor:    1921.6800.5080.00 Metric:     10
  IS neighbor:    1921.6800.5081.02 Metric:      0
  IP prefix:       192.168.8.197/32 Metric:     10 Internal
  IP prefix:       192.168.8.193/32 Metric:     10 Internal
```

Sample Output: show isis database extensive

```
user@host> show isis database extensive
IS-IS level 1 link-state database:

1921.6800.5081.00-00  Sequence: 0x89, Checksum: 0x4b4b, Lifetime: 1199 secs

  IS neighbor:    1921.6800.5080.00 Metric:     10
  IS neighbor:    1921.6800.5081.02 Metric:      0

  IP prefix:       192.168.8.197/32 Metric:     10 Internal
  IP prefix:       192.168.8.193/32 Metric:     10 Internal

  Header: LSP id: 1921.6800.5081.00-00, Length: 206 bytes
    Allocated length: 1492 bytes
    Remaining lifetime: 854 secs, Level: 2, Interface: 0
    Estimated free bytes: 1286, Actual free bytes: 1286
    Aging timer expires in: 854 secs
    Protocols: IP

  Packet: LSP id: 1921.6800.5081.00-00, Length: 206 bytes, Lifetime: 1199 secs
    Checksum: 0x4b4b, Sequence: 0x89, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 47.0005.80ff.f800.0000.0108.0001 (13)
    Speaks: IP
    IP router id: 192.168.5.81
    IP address: 192.168.5.81
    IS neighbor: 1921.6800.5081.02, Metric: 0
    IS neighbor: 1921.6800.5080.00, Metric: 10
      Maximum bandwidth: 622Mbps
      IP address: 192.168.8.192
      Neighbor's IP address: 192.168.8.193
      IP prefix: 192.168.8.193/32
      Administrative groups: 0 <none>
    IS neighbor: 1921.6800.5080.00, Metric: 10
      Maximum bandwidth: 155Mbps
      IP address: 192.168.8.196
      Neighbor's IP address: 192.168.8.197
      IP prefix: 192.168.8.197/32
      Administrative groups: 0 <none>
    IP prefix: 192.168.8.193/32, Internal, Metrics: default 10
    IP prefix: 192.168.8.197/32, Internal, Metrics: default 10
    IS neighbor: 1921.6800.5080.00, Internal, Metrics: default 10
    IS neighbor: 1921.6800.5080.00, Internal, Metrics: default 10

IS-IS level 2 link-state database:
...
```

show isis hostname

Syntax show isis hostname

Description Display hostname database.

Required Privilege Level view

Output Fields System Id—System identifier mapped to the hostname.

Hostname—Hostname mapped to the system identifier.

Type—Type of mapping between system identifier and hostname.

- Dynamic—Hostname mapping determined as described in RFC 2763.
- Static—Hostname mapping configured by user.

Sample Output

```
user@host> show isis hostname
   IS-IS hostname database:
   System Id      Hostname                                          Type
   1921.6800.4201 isis1                                             Dynamic
   1921.6800.4202 isis2                                             Static
   1921.6800.4203 isis3                                             Dynamic
```

show isis interface

Syntax show isis interface <*interface-name*> <brief | detail | extensive>

Description Display status information about the interfaces on which IS-IS is configured.

Options brief—(Optional) Display brief information about IS-IS interfaces.

detail—(Optional) Display detailed information about IS-IS interfaces.

extensive—(Optional) Display extensive information about IS-IS interfaces.

interface-name—(Optional) Name of an interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show isis interface (Standard) on page 585
Sample Output: show isis interface brief on page 586
Sample Output: show isis interface detail on page 586

Options at a Glance Table 70 summarizes which information is included in each show isis interface command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 70: Show IS-IS Interface Output Field Summary

Options	Field Description
Detail	Adjacencies—Number of adjacencies established on this interface.
Detail	Designated router—A router selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.
Detail	Circuit type—Circuit type.
All	CirID or Circuit id—Circuit identifier.
Detail	Hello (s)—Interface's hello interval.
Detail	Hold (s)—Interface's hold time.
Detail	Index—Interface index assigned by the JUNOS kernel.
Brief	Interface—Interface through which the adjacency is made.
Detail	*interface-name*—Name of the interface running IS-IS.
All	L or Level—Level.
Standard	L1/L2 Metric—Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.
Brief	Level 1 DR—Level 1 designated intermediate system.
Brief	Level 2 DR—Level 2 designated intermediate system.
Detail	LSP interval—Interface's LSP interval.
Detail	Metric—Metric value for this interface.
Detail	Priority—Priority value for this interface.
Detail	State—Internal implementation information.
Detail	Sysid—System identifier.

Output Fields

interface-name—(Detail output only) Name of the interface.

Index—(Detail output only) Interface index assigned by the JUNOS kernel.

State—(Detail output only) Internal implementation information.

Circuit id—(Detail output only) Circuit identifier.

Circuit type—(Detail output only) Circuit type:

- 1—Level 1 only
- 2—Level 2 only
- 3—Level 1 and Level 2

LSP interval—(Detail output only) Interface's LSP interval.

Sysid—(Detail output only) System identifier.

Interface—(Brief output only) Interface through which the adjacency is made.

L or Level—Level:

- 1—Level 1 only
- 2—Level 2 only
- 3—Level 1 and Level 2

CirID—Circuit identifier.

Level 1 DR—(Brief output only) Level 1 designated intermediate system.

Level 2 DR—(Brief output only) Level 2 designated intermediate system.

L1/L2 Metric—(Standard output only) Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.

Adjacencies—(Detail output only) Number of adjacencies established on this interface.

Priority—(Detail output only) Priority value for this interface.

Metric—(Detail output only) Metric value for this interface.

Hello (s)—(Detail output only) Interface's hello interval.

Hold (s)—(Detail output only) Interface's hold time.

Sample Output: show isis interface (Standard)

```
user@host> show isis interface
IS-IS interface database:
Interface L CirID Level 1 DR      Level 2 DR    L1/L2 Metric
fxp0.0 3     0x2 Disabled        Disabled       10/10
fxp1.0 3     0x3 mpls6.03        mpls6.03       10/10
lo0.0 0      0x1 Passive         Passive         0/0
sr0.0 3      0x1 Point to point  Point to point  10/10
sr1.0 3      0x1 Point to point  Point to point  10/10
```

Sample Output: show isis interface brief

```
user@host> show isis interface brief
IS-IS interface database:
Interface  L CirID Level 1 DR        Level 2 DR
lo0.0      3   0x1 crater.01         crater.01
so-1/0/0.0 2   0x9 Disabled          crater.09
so-1/1/0.0 2   0x7 Disabled          crater.07
so-1/2/0.0 2   0x3 Disabled          crater.03
so-1/3/0.0 2   0x5 Disabled          crater.05
so-2/0/0.0 2   0x2 Disabled          crater.02
so-2/1/0.0 2   0x4 Disabled          crater.04
so-2/2/0.0 2   0x6 Disabled          crater.06
so-2/3/0.0 2   0x8 Disabled          crater.08
so-5/0/0.0 2   0xa Disabled          crater.0a
so-5/0/1.0 2   0xb Disabled          crater.0b
so-5/0/2.0 2   0xc Disabled          crater.0c
so-5/0/3.0 2   0xd Disabled          crater.0d
```

Sample Output: show isis interface detail

```
user@host> show isis interface detail
lo0.0
  Index: 3, State: 0x7, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, Sysid: crater
  Level Adjacencies Priority Metric Hello (s) Hold (s)
      1           0       64      0         9       27
      2           0       64      0         9       27

so-1/0/0.0
  Index: 5, State: 0x106, Circuit id: 0x9, Circuit type: 2
  LSP interval: 100 ms
  Level Adjacencies Priority Metric Hello (s) Hold (s)
      2           1       64     10         9       27

so-1/1/0.0
  Index: 6, State: 0x106, Circuit id: 0x7, Circuit type: 2
  LSP interval: 100 ms
  Level Adjacencies Priority Metric Hello (s) Hold (s)
      2           1       64     10         9       27
...
```

show isis routes

Syntax show isis routes

Description Display the routes in the IS-IS routing table.

Required Privilege Level view

Output Fields Current version—Number of the current version of the IS-IS routing table.

L1—Version of Level 1 SPF that was run.

L2—Version of Level 2 SPF that was run.

Prefix—Destination of the route.

L—Level:

- 1—Level 1 only
- 2—Level 2 only
- 3—Level 1 and Level 2

Version—Version of SPF that generated the route.

Metric—Metric value associated with the route.

Type—Metric type. It can be int (internal) or ext (external).

Interface—Interface to the next hop.

Via—System identifier of the next hop, displayed as a name if possible.

Sample Output

```
user@host> show isis routes
IS-IS routing table Current version: L1: 89 L2: 158
Prefix L Version Metric Type Interface Via
10.1.1.1/32 2 158 20 int so-1/0/0.0 router-d
ge-2/1/0.0 router-d
10.10.20.16/29 2 158 20 int ge-4/2/0.0 router-b
10.10.20.24/29 2 158 20 int so-1/0/0.0 router-d
ge-2/1/0.0 router-d
10.10.20.102/32 2 158 10 int ge-4/2/0.0 router-b
10.10.20.103/32 2 158 20 int so-1/0/0.0 router-d
ge-2/1/0.0 router-d
ge-4/2/0.0 router-b
10.10.20.104/32 2 158 10 int so-1/0/0.0 router-d
ge-2/1/0.0 router-d
10.10.20.248/30 2 158 20 int so-1/0/0.0 router-d
ge-2/1/0.0 router-d
ge-4/2/0.0 router-b
10.10.90.0/29 2 158 20 int so-1/0/0.0 router-d ge-2/1/0.0 router-d
...
```

show isis spf

Syntax show isis spf (brief | log | results)

Description Display information about shortest-path-first (SPF) calculations.

Options brief—Display an overview of SPF calculations.

log—Display the log of SPF calculations.

results—Display the results of SPF calculations.

Required Privilege Level view

Sample Output Sample Output: show isis spf brief on page 589
Sample Output: show isis spf log on page 589
Sample Output: show isis spf results on page 590

Options at a Glance Table 71 summarizes which information is included in each show isis spf command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 71: Show ISIS SPF Interface Output Field Summary

Options	Field Description
Log	Count—Number of times the SPF was triggered.
Log	Elapsed time—Length of time required to complete the SPF computation in seconds.
Brief Results	Interface—Interface of the next hop.
Brief Results	Metric—Metric to the node.
Brief Results	Node—Sysid of a node.
Log	Reason—Reason that the SPF computation was completed.
Brief Results	SNPA—Subnetwork point of attachment (MAC address of the next hop).
Log	Start time—Time that the SPF computation started.
Brief Results	Via—Sysid of the next hop.

Output Fields Node—Sysid of a node.

Metric—Metric to the node.

Interface—Interface of the next hop.

Via—Sysid of the next hop.

SNPA—Subnetwork point of attachment (MAC address of the next hop).

Start time—(Log output only) Time that the SPF computation started.

Elapsed time—(Log output only) Length of time required to complete the SPF computation in seconds.

Count—(Log output only) Number of times the SPF was triggered.

Reason—(Log output only) Reason that the SPF computation was completed.

Sample Output: show isis spf brief

```
user@host> show isis spf brief
IS-IS level 1 SPF results:

IS-IS level 2 SPF results:
Node                  Metric Interface Via                SNPA
badlands.00               23 so-1/2/0.0 1921.6800.5067
                             so-1/0/0.0 1921.6800.5067
                             so-1/1/0.0 1921.6800.5067
                             so-1/3/0.0 1921.6800.5067
carlsbad.00               13 so-1/2/0.0 1921.6800.5067
                             so-1/0/0.0 1921.6800.5067
                             so-1/1/0.0 1921.6800.5067
                             so-1/3/0.0 1921.6800.5067
1921.6800.5064.00         13 so-1/2/0.0 1921.6800.5067
                             so-1/0/0.0 1921.6800.5067
                             so-1/1/0.0 1921.6800.5067
                             so-1/3/0.0 1921.6800.5067
1921.6800.5067.00         10 so-1/2/0.0 1921.6800.5067
                             so-1/0/0.0 1921.6800.5067
                             so-1/1/0.0 1921.6800.5067
                             so-1/3/0.0 1921.6800.5067
1921.6800.5066.00         10 so-2/2/0.0 1921.6800.5066
                             so-2/0/0.0 1921.6800.5066
                             so-2/1/0.0 1921.6800.5066
                             so-2/3/0.0 1921.6800.5066
  5 nodes
```

Sample Output: show isis spf log

```
user@host> show isis spf log
IS-IS level 1 SPF log:
Start time          Elapsed time (secs) Count Reason
Fri Aug 14 18:52:32            0.000427     1 Reconfig
Fri Aug 14 18:52:39            0.000493     1 Updated LSP
Fri Aug 14 19:05:05            0.000463     1 Periodic SPF
Fri Aug 14 19:18:47            0.000475     1 Periodic SPF
Fri Aug 14 19:31:11            0.000461     1 Periodic SPF
...

IS-IS level 2 SPF log:
Start time          Elapsed time (secs) Count Reason
Fri Aug 14 20:46:12            0.001272     1 Updated LSP
Fri Aug 14 21:03:46            0.001256     1 Updated LSP
Fri Aug 14 21:05:36            0.001108     1 Reconfig
Fri Aug 14 21:19:07            0.001808     1 Lost adjacency
Fri Aug 14 21:19:30            0.001416     1 Updated LSP
Fri Aug 14 21:30:52            0.001390     1 Periodic SPF
...
```

Sample Output: show isis spf results

```
user@host> show isis spf results
IS-IS level 1 SPF results:

IS-IS level 2 SPF results:
Node                  Metric Interface Via                SNPA
badlands.00               23 so-1/2/0.0 1921.6800.5067
                             so-1/0/0.0 1921.6800.5067
                             so-1/1/0.0 1921.6800.5067
                             so-1/3/0.0 1921.6800.5067
                          23 10.255.245.46/32
                          33 10.101.1.12/30
                          33 10.101.1.8/30
                          33 10.101.1.4/30
                          63 10.20.1.4/32
                          63 10.20.1.3/32
                          63 10.20.1.2/32
                          63 10.20.1.1/32
...
```

show isis statistics

Syntax show isis statistics

Description Display statistics about IS-IS traffic.

Required Privilege Level view

See Also clear isis statistics on page 576

Output Fields PDU type—Protocol data unit type.

Received—Number of PDUs received since IS-IS started or since the statistics were zeroed.

Processed—Number of PDUs received less the number dropped.

Drops—Number of PDUs dropped.

Sent—Number of PDUs transmitted since IS-IS started or since the statistics were zeroed.

Rexmit—Number of PDUs retransmitted since IS-IS started or since the statistics were zeroed.

Total packets received/sent—Total number of PDUs received and transmitted since IS-IS started or since the statistics were zeroed.

SNP queue length—Number of CSNP and PSNP packets currently sitting on the SNP queue waiting for processing. This value is almost always 0.

LSP queue length—Number of LSPs sitting on the LSP queue waiting for processing. This value is almost always 0.

SPF runs—Number of SPF calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.

Fragments rebuilt—Number of LSP fragments that the local system has computed.

LSP regenerations—Number of LSPs that have been regenerated. An LSP is regenerated when it is nearing the end of its lifetime and it has not changed.

Purges initiated—Number of purges that the system initiated. A purge is initiated if the software decides that an LSP must be removed from the network.

Sample Output

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12227      12227          0       8184        683
IIH            113808     113808          0     115817          0
CSNP           198868     198868          0     198934          0
PSNP             6985       6979          6       8274          0
Unknown             0          0          0          0          0
Totals         331888     331882          6     331209        683

Total packets received: 331888 Sent: 331892

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1014
Fragments rebuilt:       1038
LSP regenerations:        425
Purges initiated:           0
```

Chapter 36
OSPF Monitoring and Troubleshooting

Table 72 summarizes the command-line interface (CLI) commands you can use to monitor OSPF. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 72: Commands for Monitoring OSPF

Task or Information to Monitor	Command
Status of interfaces on which OSPF is running	show ospf interface on page 602
Link-state database entries	show ospf database on page 596
	clear ospf database on page 594
OSPF routing table entries	show ospf route on page 610
Adjacent routers	show ospf neighbor on page 607
	clear ospf neighbor on page 594
SPF log	show ospf log on page 606
OSPF input and output statistics.	show ospf io-statistics on page 605
General OSPF statistics	show ospf statistics on page 612
	clear ospf statistics on page 595

clear ospf database

Syntax clear ospf database <purge>

Description Delete the entries in the OSPF link-state advertisement database.

This command is useful only for testing purposes. Use it with care because it causes significant network perturbation.

Options none—Discard all link-state advertisements other than the system's own (the system's own link-state advertisements (LSAs) are regenerated). To resynchronize the database, all adjacent neighbors that are in the state EXSTART or higher are destroyed. They will be reacquired and the databases will be synchronized again.

purge—(Optional) Discard all entries in the link-state advertisement database; that is, all link-state advertisements are set to MAXAGE and are flooded. The database will be repopulated when the originators of the link-state advertisements receive the MAXAGE link-state advertisements and reissue them.

Required Privilege Level clear

clear ospf neighbor

Syntax clear ospf neighbor <*address*>

Description Tear down neighbor connections.

Options none—Tear down the connections with all neighbors.

address—(Optional) Tear down the connection with the specified neighbor.

Required Privilege Level clear

Output Fields Address—Address of the neighbor.

clear ospf statistics

Syntax clear ospf statistics

Description Clear OSPF statistics.

Required Privilege Level clear

See Also show ospf statistics on page 612

Sample Output

```
user@host> show ospf statistics

Packet type             Total                  Last 5 seconds
                    Sent      Received        Sent      Received
   Hello            3254          2268           3             1
     DbD              41            46           0             0
   LSReq               8             7           0             0
LSUpdate             212           154           0             0
   LSAck              65            98           0             0

LSAs retransmitted: 3, last 5 seconds: 0

Flood queue depth: 0
Total rexmit entries: 0, db summaries: 0, lsreq entries: 0

Receive errors:
  626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

Packet type             Total                  Last 5 seconds
                    Sent      Received        Sent      Received
   Hello               3             1           3             1
     DbD               0             0           0             0
   LSReq               0             0           0             0
LSUpdate               0             0           0             0
   LSAck               0             0           0             0

LSAs retransmitted: 0, last 5 seconds: 0

Flood queue depth: 0
Total rexmit entries: 0, db summaries: 0, lsreq entries: 0

Receive errors:
  None
```

show ospf database

Syntax show ospf database <summary | brief | detail | extensive> <instance *instance-name*> <*lsa-filters*>

Description Display the entries in the OSPF link-state database, which contains data about LSA packets.

Options brief—(Optional) Display a brief listing of all entries in the OSPF link-state database.

detail—(Optional) Display detailed information about the entries in the OSPF link-state database.

extensive—(Optional) Display extremely detailed information about the entries in the OSPF link-state database.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

lsa-filters—(Optional) One or more of the following LSA filters. If you specify more than one filter, only LSAs that match all the filters are displayed. For example, the command show ospf database detail router lsa-id 10.0.0.1 displays all router LSAs in all areas that have an LSA identifier of 10.0.0.1.

- advertising-router *address*—(Optional) Display the LSAs advertised by a particular router.
- area *area-id*—(Optional) Display the LSAs in a particular area.
- lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.
- *lsa-type*—(Optional) Display specific types of LSAs. You can specify asbrsummary, extern, netsummary, network, nssa, or router.

summary—(Optional) Display summary information about OSPF link-state database.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ospf database brief on page 599
Sample Output: show ospf database detail on page 600
Sample Output: show ospf database extensive on page 600
Sample Output: show ospf database summary on page 601

Options at a Glance Table 73 summarizes which information is included in each show ospf database command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 73: Show OSPF Database Output Field Summary

Options	Field Description
Brief Detail Extensive	Adv Rtr—Address of the router that sent the advertisement.
Brief Detail Extensive	Age—Time elapsed since the LSA was originated, in seconds.
Extensive	Aging timer—How long until the LSA will expire in Min:Sec.
All	area—Area number.
Brief Detail Extensive	Cksum—Checksum value of the LSA.
Extensive	expires—How long until the route expires in Hours:Min:Sec.
Summary	Externals—Display of the external link-state advertisement database.
Brief Detail Extensive	ID—Link identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local router.
Extensive	Installed—How long ago the route was installed.
Brief Detail Extensive	Len—Length of the advertisement, in bytes.
Summary	Network LSAs—Number of network link-state advertisements in the link-state database.
Extensive	Ours—Indicates that this is a local advertisement.
Summary	Router LSAs—Number of router link-state advertisements in the link-state database.
Brief Detail Extensive	Seq—Link sequence number of the advertisement.
Summary	Summary LSAs—Number of summary link-state advertisements in the link-state database.
Brief Detail Extensive	Type—Type of link advertisement. It can be ASBRSum, External, Network, NSSA, Router, or Summary.
Router Link-state Advertisements	
Detail Extensive	bits—Flags describing the router that generated the LSP.
Detail Extensive	link count—Number of links in the advertisement.
Each Link	
Detail Extensive	data—For stub networks, the subnet mask. Otherwise, the IP address of the router that generated the LSP.
Detail Extensive	id—ID of a router or subnet on the link.
Detail Extensive	TOS 0 metric—Metric for ToS 0.
Detail Extensive	TOS count—Number of type-of-service (ToS) entries in the advertisement.
Detail Extensive	type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual.
Each ToS entry	
Detail Extensive	metric—Metric for the ToS.
Detail Extensive	TOS—Type of service value.
Network Link-state Advertisements	
Detail Extensive	attached router—ID of the attached neighbor.
Detail Extensive	mask—Network mask.
Summary Link-state Advertisements	
Detail Extensive	mask—Network mask.
Detail Extensive	metric—Metric for the ToS.
Detail Extensive	TOS—Type of service value.

Output Fields area—Area number. Area 0.0.0.0 is the backbone area.

Type—Type of link advertisement. It can be ASBRSum, Extern, Network, NSSA, Router, or Summary.

ID—Link identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local router.

Adv Rtr—Address of the router that sent the advertisement.

Seq—Link sequence number of the advertisement.

Age—Time elapsed since the LSA was originated, in seconds.

Cksum—Checksum value of the LSA.

Len—Length of the advertisement, in bytes.

For router link-state advertisements:

- bits—Flags describing the router that generated the LSP.
- link count—Number of links in the advertisement.

For each link:

- id—ID of a router or subnet on the link.
- data—For stub networks, the subnet mask. Otherwise, the IP address of the router that generated the LSP.
- type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual.
- TOS count—Number of type-of-service (ToS) entries in the advertisement.
- TOS 0 metric—Metric for ToS 0.

For each ToS entry:

- TOS—Type of service (ToS) value.
- metric—Metric for the ToS.

Aging timer—(Extensive output only) How long until the LSA will expire in Min:Sec.

Installed—(Extensive output only) How long ago the route was installed.

expires—(Extensive output only) How long until the route expires in Hours:Min:Sec.

Ours—(Extensive output only) Indicates that this is a local advertisement.

For network link-state advertisements:

- mask—Network mask.
- attached router—ID of the attached neighbor.

For summary link-state advertisements:

- mask—Network mask.
- TOS—Type of service (ToS) value.
- metric—Metric for the ToS.

Router LSAs—(Summary output only) Number of router link-state advertisements in the link-state database.

Network LSAs—(Summary output only) Number of network link-state advertisements in the link-state database.

Summary LSAs—(Summary output only) Number of summary link-state advertisements in the link-state database.

Externals—(Summary output only) Display of the external link-state advertisement database.

Sample Output: show ospf database brief

```
user@host> show ospf database brief
    OSPF link state database, area 0.0.0.0
 Type       ID               Adv Rtr             Seq           Age    Cksum
Len
Router   10.250.240.8     10.250.240.8     0x800001fc  2388  0x3684  36
Router   10.250.240.17    10.250.240.17    0x80000217  1835  0x444c  36
Router   10.250.240.32    10.250.240.32    0x80000232  1876  0x0158  36
Router   10.250.240.35    10.250.240.35    0x80000291  1100  0x4aa5  36
Network  192.168.254.230  10.250.240.8     0x800001cc   117  0xab67  40
Summary  10.1.2.0         10.250.240.17    0x80000216  1535  0x1729  28
Summary  10.1.3.34        10.250.240.8     0x8000013a  2217  0x842f  28
Summary  10.1.3.34        10.250.240.35    0x800001a3   800  0x0f20  28
Summary  10.1.3.35        10.250.240.8     0x8000013b   288  0x7839  28
Summary  10.1.3.36        10.250.240.35    0x800001a4  1430  0xf833  28
Summary  10.1.2.48        10.250.240.8     0x800001f7  1317  0xaa8d  28
Summary  10.1.2.48        10.250.240.32    0x80000232   430  0xa242  28

    OSPF link state database, area 1.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Cksum  Len
Router   10.250.240.9     10.250.240.9     0x80000267   116  0x1bb3  36
Router   10.250.240.10    10.250.240.10    0x800001ce  2238  0xc6fe  36
Router   10.250.240.11    10.250.240.11    0x80000207  1896  0xca60  36
Router   10.250.240.12    10.250.240.12    0x800002a7   117  0x6029  36
Router   10.250.240.17    10.250.240.17    0x80000232   635  0x2fc7  60
Network  10.1.2.2         10.250.240.17    0x80000213  2043  0x43f5  32
Network  10.1.2.82        10.250.240.17    0x80000210   335  0x34b6  32
Network  10.1.1.130       10.250.240.17    0x800001e2   543  0x490a  36
Summary  10.1.3.34        10.250.240.17    0x8000014f  1743  0x2e66  28
Summary  10.1.3.35        10.250.240.9     0x80000168  1316  0x186b  28
Summary  10.1.3.35        10.250.240.17    0x80000150    35  0x2270  28
Summary  10.1.3.36        10.250.240.9     0x80000167  1916  0x1073  28
Summary  10.1.3.36        10.250.240.17    0x8000014f  2135  0x1a78  28
Summary  10.1.2.48        10.250.240.10    0x800001ca  1638  0x035f  28
Summary  10.1.3.64        10.250.240.10    0x800001ca   738  0x364b  28
Summary  10.1.3.64        10.250.240.17    0x8000020f   843  0x80b4  28
Summary  192.168.254.224  10.250.240.17    0x80000216  1443  0x4cb7  28

    OSPF external link state database
 Type       ID               Adv Rtr           Seq      Age  Cksum  Len
```

Sample Output: show ospf database detail

```
user@host> show ospf database detail
    OSPF link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Cksum  Len
Router   10.250.240.8     10.250.240.8     0x800001fd   464  0x3485  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.230, type Transit (2)
  TOS count 0, TOS 0 metric 1
Router   10.250.240.17    10.250.240.17    0x80000218   119  0x424d  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.227, type Transit (2)
  TOS count 0, TOS 0 metric 1
Router   10.250.240.32    10.250.240.32    0x80000233   306  0xfe59  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.225, type Transit (2)
  TOS count 0, TOS 0 metric 1
Router   10.250.240.35    10.250.240.35    0x80000291  1576  0x4aa5  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.229, type Transit (2)
  TOS count 0, TOS 0 metric 1
Network  192.168.254.230  10.250.240.8     0x800001cc   593  0xab67  40
  mask 255.255.255.240
  attached router 10.250.240.8
  attached router 10.250.240.32
  attached router 10.250.240.17
  attached router 10.250.240.35
Summary  10.1.2.0         10.250.240.17    0x80000216  2011  0x1729  28
  mask 255.255.255.240
  TOS 0x0, metric 1
Summary  10.1.3.34        10.250.240.8     0x8000013b   293  0x8230  28
  mask 255.255.255.255
  TOS 0x0, metric 65
...
```

Sample Output: show ospf database extensive

```
user@host> show ospf database extensive
    OSPF link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq      Age  Cksum  Len
Router   10.250.240.8     10.250.240.8     0x800001fd   766  0x3485  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.230, type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:47:13
  Installed 00:12:44 ago, expires in 00:47:14

Router   10.250.240.17    10.250.240.17    0x80000218   421  0x424d  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.227, type Transit (2)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:42:58  Aging timer 00:52:58
  Installed 00:07:01 ago, expires in 00:52:59
  Ours
Router   10.250.240.32    10.250.240.32    0x80000233   608  0xfe59  36
  bits 0x1, link count 1
  id 192.168.254.230, data 192.168.254.225, type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:49:51
  Installed 00:10:04 ago, expires in 00:49:52
```

```
Network  192.168.254.230  10.250.240.8     0x800001cc   895  0xab67  40
  mask 255.255.255.240
  attached router 10.250.240.8
  attached router 10.250.240.32
  attached router 10.250.240.17
  attached router 10.250.240.35
  Aging timer 00:45:05
  Installed 00:14:53 ago, expires in 00:45:05

Summary  10.1.2.0         10.250.240.17    0x80000217   121  0x152a  28
  mask 255.255.255.240
  TOS 0x0, metric 1
  Gen timer 00:47:58  Aging timer 00:57:58
  Installed 00:02:01 ago, expires in 00:57:59
  Ours
Summary  10.1.3.34        10.250.240.8     0x8000013b   595  0x8230  28
  mask 255.255.255.255
  TOS 0x0, metric 65
  Aging timer 00:50:05
  Installed 00:09:53 ago, expires in 00:50:05
...
```

Sample Output: show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.0:
   4 Router LSAs
   1 Network LSA
   13 Summary LSAs

Area 1.0.0.0:
   5 Router LSAs
   3 Network LSAs
   12 Summary LSAs

Externals:
```

show ospf interface

Syntax show ospf interface <*interface-name*> <instance *instance-name*> <brief | detail | extensive>

Description Display the status of interfaces that are running OSPF.

Options brief—(Optional) Display brief status information about OSPF interfaces.

detail—(Optional) Display detailed status information about OSPF interfaces.

extensive—(Optional) Display extremely detailed status information about OSPF interfaces, including information about what is queued on the interfaces and neighbors. When OSPF is idle, the output of the extensive option is identical to that of the detail option.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

interface-name—(Optional) Name of an interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ospf interface brief on page 604
Sample Output: show ospf interface detail on page 604
Sample Output: show ospf interface extensive on page 604

Options at a Glance Table 74 summarizes which information is included in each show ospf interface command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 74: Show OSPF Interface Output Field Summary

Options	Field Description
Extensive	Ack list—Acknowledgment list. List of pending acknowledgments on this interface.
Detail Extensive	address—IP address of the neighbor.
Detail Extensive	adj count—Number of adjacent neighbors.
All	Area—Number of the area that the interface is in.
Detail Extensive	BDR ID—Backup Designated Route for a particular subnet.
All	BDR ID—Address of the area's border router.
Detail Extensive	cost—Interface's cost (metric).
Detail Extensive	Dead—Dead timer.
Extensive	Descriptor list—List of packet descriptors.
Detail Extensive	DR addr—Address of the designated router.
All	DR ID—Address of the area's designated router.
Extensive	Flood list—List of LSAs pending flood on this interface.
Detail Extensive	Hello—Hello timer.

Options	Field Description
All	Interface—Name of the interface running OSPF.
Detail Extensive	mask—Mask of the interface.
Detail Extensive	MTU—Interface's MTU.
All	Nbrs—Number of neighbors on this interface.
Detail Extensive	OSPF area type—Time of OSPF area.
Detail	Priority—The router priority used in DR election on this interface.
Detail Extensive	ReXmit—Retransmit timer.
All	State—State of the interface. It can be BDR, Down, DR, DRother, Loop, PtToPt, or Waiting.
Detail Extensive	Type—Type of interface.

Output Fields

Interface—Name of the interface running OSPF.

State—State of the interface. It can be BDR, Down, DR, DRother, Loop, PtToPt, or Waiting.

Area—Number of the area that the interface is in.

DR ID—Address of the area's designated router.

BDR ID—Backup Designated Route for a particular subnet.

Nbrs—Number of neighbors on this interface.

Type—(Detail and Extensive output only) Type of interface. It can be LAN, NBMA, P2MP, P2P, or Virtual.

address—(Detail and Extensive output only) IP address of the neighbor.

mask—(Detail and Extensive output only) Mask of the interface.

MTU—(Detail and Extensive output only) Interface's MTU.

cost—(Detail and Extensive output only) Interface's cost (metric).

DR addr—(Detail and Extensive output only) Address of the designated router.

BDR addr—Address of the backup designated router.

adj count—(Detail and Extensive output only) Number of adjacent neighbors.

Priority—The router priority used in designated router (DR) election on this interface.

Flood list—(Extensive output only) List of LSAs pending flood on this interface.

Ack list—(Extensive output only) Acknowledgment list. List of pending acknowledgments on this interface.

Descriptor list—(Extensive output only) List of packet descriptors.

Hello—(Detail and Extensive output only) Configured value for the Hello timer.

Dead—(Detail and Extensive output only) Configured value for the Dead timer.

ReXmit—(Detail and Extensive output only) Configured value for the Retransmit timer.

OSPF area type—(Detail and Extensive output only) Type of OSPF area. It can be Stub, Not Stub, or Stub NSSA.

Sample Output: show ospf interface brief

```
user@host> show ospf interface brief
Intf                State     Area            DR ID           BDR ID       Nbrs
at-5/1/0.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
ge-2/3/0.0          DR       0.0.0.0         192.168.4.16    192.168.4.15    1
lo0.0               DR       0.0.0.0         192.168.4.16    0.0.0.0         0
so-0/0/0.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/1.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
so-6/0/2.0          Down     0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/3.0          PtToPt   0.0.0.0         0.0.0.0         0.0.0.0         1
```

Sample Output: show ospf interface detail

```
user@host> show ospf interface detail
Interface              State     Area            DR ID           BDR ID
Nbrs
fe-0/0/1.0             BDR      0.0.0.0         192.168.37.12   10.255.245.215
1
Type LAN, address 192.168.37.11, mask 255.255.255.248, MTU 4460, cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, adj count 1, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
lo0.0                  DR       0.0.0.0         10.255.245.215  0.0.0.0
0
Type LAN, address 10.255.245.215, mask 255.255.255.255, MTU 65535, cost 0
DR addr 10.255.245.215, BDR addr 0.0.0.0, adj count 0, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
lo0.0                  DR       0.0.0.0         10.255.245.215  0.0.0.0
0
Type LAN, address 127.0.0.1, mask 255.255.255.255, MTU 65535, cost 0
DR addr 127.0.0.1, BDR addr 0.0.0.0, adj count 0, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
```

Sample Output: show ospf interface extensive

```
user@host> show ospf interface extensive
Interface              State     Area            DR ID           BDR ID
Nbrs
fe-0/0/1.0             BDR      0.0.0.0         192.168.37.12   10.255.245.215
1
Type LAN, address 192.168.37.11, mask 255.255.255.248, MTU 4460, cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, adj count 1, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
lo0.0                  DR       0.0.0.0         10.255.245.215  0.0.0.0
0
Type LAN, address 10.255.245.215, mask 255.255.255.255, MTU 65535, cost 0
DR addr 10.255.245.215, BDR addr 0.0.0.0, adj count 0, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
lo0.0                  DR       0.0.0.0         10.255.245.215  0.0.0.0
0
Type LAN, address 127.0.0.1, mask 255.255.255.255, MTU 65535, cost 0
DR addr 127.0.0.1, BDR addr 0.0.0.0, adj count 0, priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
```

show ospf io-statistics

Syntax show ospf io-statistics <instance *instance-name*>

Options: instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

Description Display OSPF input and output statistics.

Required Privilege Level view

Output Fields Packets read—Number of OSPF packets read.

average per run—OSPF read is a process that periodically occurs since the last time the routing protocol daemon was started. The total number of packets divided by the total number of times OSPF read is scheduled is the average per run.

max run—Maximum number of packets for a given run among all scheduled runs.

Received errors—Number of faulty packets received with errors.

Sample Output: show ospf io-statistics

```
user@host> show ospf io-statistics
Packets read: 7361, average per run: 1.00, max run: 1

Receive errors:
  None
```

show ospf log

Syntax show ospf log <instance *instance-name*>

Description Display the entries in the OSPF log of shortest-path-first (SPF) calculations.

Required Privilege Level view

Options instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

Output Fields When—Time when the SPF calculation was made.

Type—Type of calculation. It can be Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or Virtuallink.

Elapsed—How much time has passed since the calculation was made, in seconds.

Sample Output

```
user@host> show ospf log
When             Type            Elapsed

1w4d 17:25:58   Stub            0.000017
1w4d 17:25:58   SPF             0.000070
1w4d 17:25:58   Stub            0.000019
1w4d 17:25:58   Interarea       0.000054
1w4d 17:25:58   External        0.000005
1w4d 17:25:58   Cleanup         0.000203
1w4d 17:25:58     Total         0.000537
1w4d 17:24:48   SPF             0.000125
1w4d 17:24:48   Stub            0.000017
1w4d 17:24:48   SPF             0.000100
1w4d 17:24:48   Stub            0.000016
1w4d 17:24:48   Interarea       0.000056
1w4d 17:24:48   External        0.000005
1w4d 17:24:48   Cleanup         0.000238
1w4d 17:24:48     Total         0.000600
...
```

show ospf neighbor

Syntax show ospf neighbor <instance *instance-name*> < neighbor *neighbor-address* > <brief | detail | extensive>

Description Display information about all OSPF neighbors.

Options brief—(Optional) Display brief information about OSPF neighbors.

detail—(Optional) Display detailed information about OSPF neighbors.

extensive—(Optional) Display extremely detailed information about OSPF neighbors.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

neighbor *neighbor-address*—(Optional) Address or ID of specific OSPF neighbor.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ospf neighbor brief on page 608
Sample Output: show ospf neighbor detail on page 608
Sample Output: show ospf neighbor extensive on page 609

Options at a Glance Table 75 summarizes which information is included in each show ospf neighbor command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 75: Show OSPF Neighbor Output Field Summary

Options	Field Description
All	Address—Address of the neighbor.
Detail Extensive	adjacent—Length of time since the adjacency with the neighbor was established.
Detail Extensive	area—Area that the neighbor is in.
Detail Extensive	BDR—Address of the backup designated router.
All	Dead—Number of seconds until the neighbor becomes unreachable.
Detail Extensive	DR—Address of the designated router.
All	ID—Router ID of the neighbor.
All	Intf—Interface through which the neighbor is reachable.
Detail Extensive	opt—Option bits from the neighbor.
All	Pri—Priority of the neighbor to become the designated router.
All	State—State of the neighbor. It can be Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2Way.
Detail Extensive	Up—Length of time since the neighbor came up.

Output Fields

Address—Address of the neighbor.

Intf—Interface through which the neighbor is reachable.

State—State of the neighbor. It can be Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2Way.

ID—Router ID of the neighbor.

Pri—Priority of the neighbor to become the designated router.

Dead—Number of seconds until the neighbor becomes unreachable.

area—(Detail and Extensive output only) Area that the neighbor is in.

opt—(Detail and Extensive output only) Option bits from the neighbor.

DR—(Detail and Extensive output only) Address of the designated router.

BDR—(Detail and Extensive output only) Address of the backup designated router.

Up—(Detail and Extensive output only) Length of time since the neighbor came up.

adjacent—(Detail and Extensive output only) Length of time since the adjacency with the neighbor was established.

Sample Output: show ospf neighbor brief

```
user@host> show ospf neighbor brief
  Address          Intf               State      ID              Pri  Dead
192.168.254.225  fxp3.0               2Way      10.250.240.32    128   36
192.168.254.230  fxp3.0               Full      10.250.240.8     128   38
192.168.254.229  fxp3.0               Full      10.250.240.35    128   33
10.1.1.129       fxp2.0               Full      10.250.240.12    128   37
10.1.1.131       fxp2.0               Full      10.250.240.11    128   38
10.1.2.1         fxp1.0               Full      10.250.240.9     128   32
10.1.2.81        fxp0.0               Full      10.250.240.10    128   33
```

Sample Output: show ospf neighbor detail

```
user@host> show ospf neighbor detail
  Address          Intf               State      ID              Pri  Dead
192.168.254.225  fxp3.0               2Way      10.250.240.32    128   35
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 18:09:18
192.168.254.230  fxp3.0               Full      10.250.240.8     128   32
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 19:21:39, adjacent 1w5d 19:21:38
192.168.254.229  fxp3.0               Full      10.250.240.35    128   31
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 19:21:39, adjacent 1w5d 19:21:39
10.1.1.129       fxp2.0               Full      10.250.240.12    128   30
  area 1.0.0.0, opt 0x2, DR 10.1.1.130, BDR 10.1.1.131
  Up 1w1d 18:09:20, adjacent 1w1d 18:09:20
10.1.1.131       fxp2.0               Full      10.250.240.11    128   37
  area 1.0.0.0, opt 0x2, DR 10.1.1.130, BDR 10.1.1.131
  Up 1w4d 18:15:02, adjacent 1w4d 18:15:02
10.1.2.1         fxp1.0               Full      10.250.240.9     128   32
  area 1.0.0.0, opt 0x2, DR 10.1.2.2, BDR 10.1.2.1
  Up 1w5d 19:21:38, adjacent 1w5d 19:20:58
10.1.2.81        fxp0.0               Full      10.250.240.10    128   32
  area 1.0.0.0, opt 0x2, DR 10.1.2.82, BDR 10.1.2.81
  Up 1w5d 18:34:57, adjacent 1w5d 18:34:19
```

Sample Output: show ospf neighbor extensive

```
user@host> show ospf neighbor extensive
  Address          Intf               State      ID               Pri  Dead
192.168.254.225  fxp3.0               2Way      10.250.240.32    128   35
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 18:09:18
192.168.254.230  fxp3.0               Full      10.250.240.8     128   32
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 19:21:39, adjacent 1w5d 19:21:38
192.168.254.229  fxp3.0               Full      10.250.240.35    128   31
  area 0.0.0.0, opt 0x2, DR 192.168.254.230, BDR 192.168.254.229
  Up 1w5d 19:21:39, adjacent 1w5d 19:21:39
10.1.1.129       fxp2.0               Full      10.250.240.12    128   30
  area 1.0.0.0, opt 0x2, DR 10.1.1.130, BDR 10.1.1.131
  Up 1w1d 18:09:20, adjacent 1w1d 18:09:20
10.1.1.131       fxp2.0               Full      10.250.240.11    128   37
  area 1.0.0.0, opt 0x2, DR 10.1.1.130, BDR 10.1.1.131
  Up 1w4d 18:15:02, adjacent 1w4d 18:15:02
10.1.2.1         fxp1.0               Full      10.250.240.9     128   32
  area 1.0.0.0, opt 0x2, DR 10.1.2.2, BDR 10.1.2.1
  Up 1w5d 19:21:38, adjacent 1w5d 19:20:58
10.1.2.81        fxp0.0               Full      10.250.240.10    128   32
  area 1.0.0.0, opt 0x2, DR 10.1.2.82, BDR 10.1.2.81
  Up 1w5d 18:34:57, adjacent 1w5d 18:34:19
```

show ospf route

Syntax show ospf route <detail> <instance *instance-name*> <(abr | asbr | extern | inter | intra)>

Description Display the entries in the OSPF routing table.

Options none—Display all entries in the OSPF routing table.

abr—(Optional) Display OSPF routes to area border routers.

asbr—(Optional) Display OSPF routes to autonomous system border routers.

detail—(Optional) Display detailed information about the entries in the OSPF routing table.

extern—(Optional) Display external OSPF routes.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

inter—(Optional) Display interarea OSPF routes.

intra—(Optional) Display intra-area OSPF routes.

Required Privilege Level view

Sample Output Sample Output: show ospf route on page 611
Sample Output: show ospf route detail on page 612

Options at a Glance Table 76 summarizes which information is included in each show ospf route command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 76: Show OSPF Route Output Fields Summary

Options	Field Description
Detail	area—Area ID of the route.
All	Metric—Route's metric value.
All	Next hop addr—Address of the next hop.
All	Next hop i/f—Interface through which the route's next hop is reachable.
Detail	optional-capability—Optional capabilities propagated as VEB bits of the router LSA of a router.
Detail	origin—Router from which the route was learned.
All	Prefix—Destination of the route.
All	Route/Path Type—How the route was learned.

Output Fields

Prefix—Destination of the route.

Route/Path Type—How the route was learned:

- ABR—Route to area border router
- AS BR—Route to AS border router
- Ext—External router
- Inter—Interarea route
- Intra—Intra-area route
- Network—Network router

Metric—Route's metric value.

Next hop i/f—Interface through which the route's next hop is reachable.

Next hop addr—Address of the next hop.

area—(Detail output only) Area ID of the route.

optional-capability—(Detail output only) Optional capabilities propagated in the router LSA. This field is in the output for router routes only, not for network routes. Currently, three bits in this field are defined: V, E, and B bits. If V is set, the router is at the end of a virtual active link. If E is set, it is an Autonomous System boundary router. If B is set, it is an area border router.

origin—(Detail output only) Router from which the route was learned.

Sample Output: show ospf route

```
user@host> show ospf route
Prefix               Route/Path Type      Metric  Next hop i/f  Next hop addr
1.1.1.0/24           Ext2    Network      0       fxp0.0        10.10.0.16
1.1.1.1/32           Intra   AS BR        1       fxp0.0        10.10.0.16
1.2.3.0/24           Ext2    Network      0       fxp0.0        10.10.0.16
10.10.0.0/24         Intra   Network      1       fxp0.0
10.10.1.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
10.10.2.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
10.10.3.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
192.168.0.0/23       Ext2    Network      0       fxp0.0        10.10.0.16
```

Sample Output: show ospf route detail

```
user@host> show ospf route detail
Prefix               Route/Path Type      Metric  Next hop i/f  Next hop addr
1.1.1.0/24           Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
1.1.1.1/32           Intra   AS BR        1       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
1.2.3.0/24           Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
10.10.0.0/24         Intra   Network      1       fxp0.0
  area 0.0.0.0, optional capability 0x0, origin 10.10.0.15
10.10.1.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
10.10.2.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
10.10.3.0/24         Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
192.168.0.0/23       Ext2    Network      0       fxp0.0        10.10.0.16
  area 0.0.0.0, optional capability 0x0, origin 1.1.1.1
```

show ospf statistics

Syntax

show ospf statistics <instance *instance-name*>

Description

Display OSPF protocol statistics.

Required Privilege Level

view

Options

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

Output Fields

Packet type—Type of OSPF packet.

Total Sent/Total Received—Total number of packets sent and received.

Last 5 seconds Sent/Last 5 seconds Received—Total number of packets sent and received in the last 5 seconds.

LSAs retransmitted—Total number of link-state advertisements transmitted, and number retransmitted in the last 5 seconds.

Receive errors—Number and type of receive errors.

Sample Output

```
user@host> show ospf statistics
Packet type             Total                  Last 5 seconds
                    Sent      Received        Sent      Received
   Hello          505739        990495           4             5
     DbD              20            26           0             0
   LSReq               6             5           0             0
LSUpdate           27060         15319           0             0
   LSAck           10923         52470           0             0

LSAs retransmitted: 16, last 5 seconds: 0

Receive errors:
  862 no interface found
  115923 no virtual link found
```

Chapter 37
BGP Monitoring and Troubleshooting

Table 77 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot BGP. In the table, the commands are grouped by functionality. In the remainder of this chapter, the commands are explained alphabetically.

Table 77: Commands for Monitoring BGP

Task or Information to Monitor	Command
Entries in the BGP group database.	show bgp group on page 615
Entries in the BGP neighbor database.	show bgp neighbor on page 617
Remove entries from the neighbor database.	clear bgp neighbor on page 614
BGP summary information.	show bgp summary on page 622
Remove damping information.	clear bgp damping on page 613

clear bgp damping

Syntax clear bgp damping <*prefix*>

Description Reset BGP damping information.

Options *prefix*—(Optional) Remove damping information for the specified destination prefix.

Required Privilege Level clear

Sample Output

```
user@host> clear bgp damping
```

clear bgp neighbor

Syntax clear bgp neighbor <*neighbor* | as *as-number*> <soft | soft-inbound>

Description Change the state of BGP neighbor to IDLE. For neighbors that are in the ESTABLISHED state, this action causes the TCP connection to that neighbor to be torn down (and subsequently re-established).

Options as *as-number*—(Optional) Clear the state of the neighbors in the specified AS only.

neighbor—(Optional) Clear the state of the specified neighbor only.

soft—(Optional) Reapply any export policies and send refresh updates to the peer without clearing the state.

soft-inbound—(Optional) Reapply any import policies and send refresh updates to the peer without clearing the state.

Required Privilege Level clear

Sample Output

```
user@host> clear bgp neighbor
```

show bgp group

Syntax show bgp group <*group-name*>

Description Display information about the configured BGP groups.

Options none—Display group information about all BGP groups.

group-name—(Optional) Display group information for the specified group.

Required Privilege Level view

Output Fields Group type—Type of BGP group. It can be either Internal or External.

AS—AS number of the peer. For IBGP, this number should be the same as Local AS.

Local AS—AS number of the local router.

Name—Name of specific BGP group.

Export—Export policies configured for the BGP group with the export statement.

Options—Configured BGP options. It can be one or more of the following:

- Local address—Address configured with the local-address statement.
- NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any.
- Hold time—Hold time configured with the hold-time statement. The default hold time is 90 seconds.
- Preference—Preference value configured with the preference statement. The default preference value is 170.

Total peers—Total number of peers in the group.

Established—Number of peers within the group that are in the established state.

ip-addresses—List of peers who are members of the group. The address is followed by the peer's port number.

Route Queue Timer—The number of seconds until queued routes will be sent out. If this time has already past, the number of seconds by which the updates are delayed.

Route Queue—Number of prefixes that are queued up for sending to the peers in the group.

Sample Output

```
user@host> show bgp group
Group Type: External                    Local AS: 65004
  Name: pe-pe
  Export: [ match-all ]
  Options: <Multihop Confed>
  Total peers: 1        Established: 0
  11.134.0.3+179

Group Type: External                    Local AS: 65004
  Name: pe-ce
  Export: [ gr-export ]
  Total peers: 1        Established: 1
  10.145.0.5+179
  Route Queue Timer: unset Route Queue: empty
```

show bgp neighbor

Syntax show bgp neighbor <*neighbor-address*>

Description Display information about BGP peers.

Options none—Display information about all BGP peers.

neighbor-address—(Optional) IP address of a BGP peer.

Required Privilege Level view

See Also For information about the local-address, nlri, hold-time, and preference statements, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols.*

Output Fields Peer—Address of the BGP neighbor. The address is followed by the neighbor's port number.

AS—AS number of the peer.

Local—Address of the local router. The address is followed by the peer's port number.

Type—Type of peer. It can be either Internal or External.

State—Current state of the BGP session. It can be one of the following:

- Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message.
- Connect—BGP is waiting for the transport protocol connection to complete.
- Established—The BGP session has been established, and the peers are exchanging update messages.
- Idle—This is the first stage of a connection. BGP is waiting for a Start event.
- OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.
- OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Flags—Internal BGP flags. It can be one or more of the following:

- CleanUp—The peer session is being shut down.
- Delete—This peer has been deleted.
- Idled—This peer has been permanently idled.
- Initializing—The peer session is initializing.
- SendRtn—Messages are being sent to the peer.
- TryConnect—Another attempt is being made to connect to the peer.

- Unconfigured—This peer is not configured.
- WriteFailed—An attempt to write to this peer failed.

Last State—Previous state of the BGP session.

Last Event—Last activity that occurred in the BGP session. It can be one of the following:

- Closed—The BGP session closed.
- ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect.
- HoldTime—The session ended because the hold timer expired.
- KeepAlive—The local router sent a BGP keepalive message to the peer.
- Open—The local router sent a BGP open message to the peer.
- OpenFail—The local router did not receive an acknowledgment of a BGP open message from the peer.
- RecvKeepAlive—The local router received a BGP keepalive message from the peer.
- RecvNotify—The local router received a BGP notification message from the peer.
- RecvOpen—The local router received a BGP open message from the peer.
- RecvUpdate—The local router received a BGP update message from the peer.
- Start—The peering session started.
- Stop—The peering session stopped.
- TransportError—A TCP error occurred.

Last Error—Last error that occurred in the BGP session. It can be one of the following:

- Cease—An error occurred, such as a version mismatch, that caused the session to close.
- Finite State Machine Error—In setting up the session, BGP received a message that it did not understand.
- Hold Time Expired—The session's hold time expired.
- Message Header Error—The header of a BGP message was malformed.
- Open Message Error—A BGP open message contained an error.
- None—No errors occured in the BGP session.
- Update Message Error—A BGP update message contained an error.

Export—Name of the export policy that is configured on the peer.

Import—Name of the import policy that is configured on the peer.

Options—Configured BGP options. It can be one or more of the following:

- AddressFamily—Configured address family. It can be inet or inet-vpn.
- GracefulRestart—Graceful restart is configured.
- HoldTime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.
- Local Address—Address configured with the local-address statement.
- NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any.
- Peer AS—Configured peer AS.
- Preference—Preference value configured with the preference statement. The default preference value is 170.
- Refresh—Configured to refresh automatically when policy changes.
- Rib-group—Configured routing table group.

Address families configured—Names of configured address families for VPN.

Local Address—The interface used by an external BGP session (except when you use the BGP multihop option for this peer).

Holdtime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.

Preference—Preference value configured with the preference statement. The default preference value is 170.

Number of flaps—Number of times the BGP session has gone down and then come back up.

Peer ID—Router identifier of the peer.

Local ID—Router identifier of the local router.

Active hold time—Hold time the local router negotiated with the peer.

Keepalive interval—Keepalive interval, in seconds.

Local Interface—Name of directly connected interface over which the direct EBGP peering is established.

NLRI for restart configured on peer—Names of address families configured for restart.

NLRI advertised by peer—Address families supported by the peer. It can be unicast or multicast.

NLRI for this session—Address families being used for this session.

Peer supports Refresh capability—Remote peer supports Route Refresh Capability, and is able to send and request full route table readvertisements. For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.

Restart time configured on the peer—Configured time allowed for restart on the neighbor. The range is 1-600 seconds. The default value is 300 seconds.

Stale routes from peer are kept for—When graceful restart is negotiated, this is the maximum time allowed to hold routes from neighbors after the BGP session has gone down. The range is 1-600 seconds. The default value is 300 seconds.

Restart time requested by this peer—Restart time requested by this neighbor during capability negotiation.

NLRI that peer supports restart for—Address families indicates graceful restart during capability negotiation.

NLRI that restart is negotiated for—Both address families are supported by both this peer and its neighbor.

NLRI of received end-of-rib markers—Address families for which end-of-rib markers are received from the neighbor.

NLRI of all end-of-rib markers sent—Address families for which end-of-rib markers are sent to the neighbor.

Table inet.*number*—Information about the routing table.

- RIB State—BGP is in the graceful restart process for this routing table. It can be restart is complete or restart in progress.
- Bit—Number that represents the entry in the routing table for this peer.
- Send state—State of the BGP group. It can be in sync, not in sync, or not advertising.
- Active prefixes—Number of prefixes received from the peer that are active in the routing table.
- Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table.
- Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

Last traffic (seconds)—Last time any traffic was received from the peer or sent to the peer, and the last time the local router checked.

Input messages—Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output messages—Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output queue[0], Output queue[1]—Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.

Trace options—Configured tracing of BGP protocol packets and operations. See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Trace file—Name of the file to receive the output of the tracing operation.

Sample Output

```
user@host> show bgp neighbor 192.168.1.111
Peer: 192.168.1.111+1067 AS 25    Local: 192.168.1.112+179 AS 26
  Type: External    State: Established    Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference HoldTime GracefulRestart AddressFamily PeerAS Refresh>
  Address families configured: inet-unicast inet6-unicast
  Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.1.111    Local ID: 192.168.1.112    Active Holdtime: 90
  Keepalive Interval: 30
  Local Interface: so-1/0/0.0
  NLRI for restart configured on peer: inet-unicast inet6-unicast
  NLRI advertised by peer: inet-unicast inet6-unicast
  NLRI for this session: inet-unicast inet6-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages:  Total 9      Updates 6       Refreshes 0     Octets 403
  Output messages: Total 7      Updates 3       Refreshes 0     Octets 365
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10
```

show bgp summary

Syntax show bgp summary

Description Display summary information about BGP and its neighbors.

Required Privilege Level view

Output Fields Groups—Number of BGP groups.

Peers—Number of BGP peers.

Down peers—Number of down BGP peers.

Table—Name of routing table.

Tot Paths—Total number of paths.

Act Paths—Number of active routes.

Suppressed—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

History—Number of withdrawn routes stored locally to keep track of damping history.

Damp State—Number of routes with a Figure of Merit (FOM) greater than zero, but still active because the FOM has not reached the threshold at which suppression occurs.

Pending—Routes in process by BGP import policy.

Peer—Address of each BGP peer. Each peer has one line of output.

AS—Peer's AS number.

InPkt—Number of packets received from the peer.

OutPkt—Number of packets sent to the peer.

OutQ—Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.

Flaps—Number of times the BGP session has gone down and then come back up.

Last Up/Down—Last up/down time since the neighbor transitioned to or from the established state.

State|#Active/Received/Damped...—Displays either the BGP state or, if the neighbor is in the established state, the number of paths received from the neighbor, as well as the number of these paths that have been accepted as active and hence are being used for forwarding.

Sample Output

```
user@host> show bgp summary
Groups: 1       Peers: 1 Down Peers: 0
Table    Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0         100        100          0          0          0          0
inet.2           0          0          0          0          0          0
  Peer             AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn State|#A
ctive/Received/Damped...
10.168.1.222        1      13433      13430       0       0 4d 15:55:10 100/100/
0            0/0/0
```

Chapter 38
RIP Monitoring and Troubleshooting

Table 78 summarizes the command-line interface (CLI) commands you can use to monitor the RIP. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 78: Commands for Monitoring RIP

Task or Information to Monitor	Command
Neighbors with which this router exchanges information.	show rip neighbor on page 626
RIP packet and error counters.	show rip statistics on page 627
	clear rip statistics on page 625

clear rip statistics

Syntax clear rip statistics <*name* | *instance*>

Description Reset RIP neighbor counters to zero.

Required Privilege Level clear

See Also show rip statistics on page 627

Options none—Reset RIP counters for all neighbors.

name—(Optional) Reset RIP counters for specified RIP neighbor.

instance—(Optional) Reset RIP counters for specified instance.

show rip neighbor

Syntax show rip neighbor <*name*>

Description Display information about all RIP neighbors.

Options none—Display brief information about all RIP neighbors.

name—(Optional) Display detailed information about a specific RIP neighbor.

Required Privilege Level view

Output Fields Neighbor—Name of RIP neighbor.

State—State of the connection. The interface can be either Up or Dn (Down).

Source Address—Source address.

Destination Address—Destination address.

Send Mode—Send options. It can be broadcast, multicast, none, or version 1.

Receive Mode—Type of packets to accept. It can be both, none, version 1, or version 2.

In Met—Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output: show rip neighbor

```
user@host> show rip neighbor
                          Source          Destination     Send   Receive   In
Neighbor          State   Address         Address         Mode   Mode     Met
--------          -----   -------         -----------     ----   -------  ---
ge-2/3/0.0           Up   192.168.9.105   192.168.9.107   bcast  both       1
at-5/1/1.42          Dn   (null)          (null)          mcast  v2 only    3
at-5/1/0.42          Dn   (null)          (null)          mcast  both       3
at-5/1/0.0           Up   20.0.0.1        224.0.0.9       mcast  both       3
so-0/0/0.0           Up   192.168.9.97    224.0.0.9       mcast  both       3
```

show rip statistics

Syntax show rip statistics <*name*>

Description Display RIP protocol statistics.

Options none—Display brief information about all RIP neighbors.

name—(Optional) Display detailed information about a specific RIP neighbor.

Required Privilege Level view

Output Fields RIP info—Information about RIP on the specified interface.

- port—UDP port number used for RIP.
- update interval—Number of seconds since last update.
- holddown—Hold-down interval, in seconds.
- timeout—Timeout interval, in seconds.
- bad msgs—Number of bad messages received.
- rts learned—Number of routes learned through RIP.
- rts held down—Number of routes held down by RIP.
- rqsts dropped—Number of received request packets that were dropped.
- resps dropped—Number of received response packets that were dropped.
- curr memory—Amount of memory currently used.
- max memory—Highwater mark of memory used by RIP.

Logical interface—Name of the logical interface.

routes learned—Number of routes learned on the logical interface.

routes advertised—Number of RIP routes advertised.

Counter—List of counter types.

Total—Total number of packets for the selected counter.

Last 5 min—Number of packets for the selected counter in the most recent five minute period.

Last minute—Number of packets for the selected counter in the most recent one minute period.

Updates Sent—Number of update messages sent.

Triggered Updates Sent—Number of triggered update messages sent.

Responses Sent—Number of response messages sent.

Bad Messages—Number of bad messages received.

RIPv1 Updates Received—Number of RIPv1 Update messages received.

RIPv1 Bad Route Entries—Number of RIPv1 Bad Route Entry messages received.

RIPv1 Updates Ignored—Number of RIPv1 Update messages ignored.

RIPv2 Updates Received—Number of RIPv2 Update messages received.

RIPv2 Bad Route Entries—Number of RIPv2 Bad Route Entry messages received.

RIPv2 Updates Ignored—Number of RIPv2 updates ignored.

Authentication Failures—Number of received update messages that failed authentication.

RIP Requests Received—Number of RIP Request messages received.

RIP Requests Ignored—Number of RIP Request messages ignored.

Sample Output: show rip statistics

```
user@host> show rip statistics so-0/0/0.0
RIP info:  port 520; update interval 30s; holddown 180s; timeout 120s.
      bad msgs  rts learned  rts held down  curr memory  max memory
             0            0              0         3840        3840

so-0/0/0.0:  0 routes learned; 5 routes advertised
Counter                          Total   Last 5 min  Last minute
-------                     -----------  -----------  -----------
Updates Sent                          0            0            0
Triggered Updates Sent                1            0            0
Responses Sent                        0            0            0
Bad Messages                          0            0            0
RIPv1 Updates Received                0            0            0
RIPv1 Bad Route Entries               0            0            0
RIPv1 Updates Ignored                 0            0            0
RIPv2 Updates Received                0            0            0
RIPv2 Bad Route Entries               0            0            0
RIPv2 Updates Ignored                 0            0            0
Authentication Failures               0            0            0
RIP Requests Received                 0            0            0
RIP Requests Ignored                  0            0            0
```

Chapter 39
RIPng Monitoring and Troubleshooting

Table 79 summarizes the command-line interface (CLI) commands you can use to monitor the RIPng. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 79: Commands for Monitoring RIPng

Task or Information to Monitor	Command
Neighbors with which this router exchanges information.	show ripng neighbor on page 630
RIPng packet and error counters.	clear ripng statistics on page 629
	show ripng statistics on page 632
Display entries in the RIPng routing tables.	show route protocol ripng table inet6 on page 631

clear ripng statistics

Syntax clear ripng statistics <*name* | *instance*>

Description Reset RIPng neighbor counters to zero.

Required Privilege Level clear

See Also show ripng statistics on page 632

Options none—Reset RIPng counters for all neighbors.

name—(Optional) Reset RIPng counters for specified RIPng neighbor.

instance—(Optional) Reset RIPng counters for specified instance.

show ripng neighbor

Syntax show ripng neighbor <*name*>

Description Display information about all RIPng neighbors.

Options none—Display brief information about all RIPng neighbors.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level view

Output Fields Neighbor—Name of RIPng neighbor.

State—State of the connection. The interface can be either Up or Dn (Down).

Source Address—Source address.

Dest Address—Destination address.

Send—Send options. It can be broadcast, multicast, none, version 1 or yes.

Receive—Type of packets to accept. It can be both, none, version 1, or yes.

In Met—Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output: show ripng neighbor

```
user@host> show ripng neighbor

                    Source                          Dest              In
Neighbor     State  Address                         Address  Send Recv Met
--------     -----  -------                         -------  ---- ---- ---
fe-0/0/2.0      Up  fe80::290:69ff:fe68:b002        ff02::9  yes yes    1
```

show route protocol ripng table inet6

Syntax show route protocol ripng table inet6

Description Display the entries in the RIPng routing tables.

Required Privilege Level view

Output Fields None.

Sample Output

```
user@host> show route protocol ripng table inet6

inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show ripng statistics

Syntax show ripng statistics <*name*>

Description Display RIPng statistics about messages sent and received on an interface, as well as information received from advertisements from other routers.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level view

See Also clear ipv6 router-advertisement on page 636

Output Fields RIPng info—Information about RIPng on the specified interface.

- port—UDP port number used for RIPng.
- update interval—Number of seconds since last update.
- holddown—Hold-down interval, in seconds.
- timeout—Timeout interval, in seconds.
- bad msgs—Number of bad messages received.
- rts learned—Number of routes learned through RIPng.
- rts held down—Number of routes held down by RIPng.
- rqsts dropped—Number of received request packets that were dropped.
- resps dropped—Number of received response packets that were dropped.
- curr memory— Amount of memory currently used.
- max memory—Highwater mark of memory used by RIPng.

Logical interface—Name of the logical interface.

routes learned—Number of routes learned on the logical interface.

routes advertised—Number of RIPng routes advertised.

Updates Sent—Number of update messages sent.

Triggered Updates Sent—Number of triggered update messages sent.

Responses Sent—Number of response messages sent.

Bad Messages—Number of bad messages received.

Updates Received—Number of RIPng update messages received.

Bad Route Entries—Number of RIPng bad route entry messages received.

Updates Ignored—Number of RIPng update messages ignored.

Authentication Failures—Number of received update messages that failed authentication.

RIPng Requests Received—Number of RIPng request messages received.

RIPng Requests Ignored—Number of RIPng request messages ignored.

Sample Output

```
user@host> show ripng statistics
RIPng info:  port 521; update interval 30s; holddown 180s; timeout 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
           4110              8              0        1646289

so-1/3/0.99:  0 routes learned; 4110 routes advertised
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                    28370          840          168
Triggered Updates Sent             42            1            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
Updates Received                 3544           30            4
Bad Route Entries                   0            0            0
Updates Ignored                     0            0            0
RIPng Requests Received             0            0            0
RIPng Requests Ignored              0            0            0
```

Chapter 40
IPv6 Monitoring and Troubleshooting

Table 80 summarizes the command-line interface (CLI) commands you can use to monitor IPv6. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 80: Commands for Monitoring IPv6

Task or Information to Monitor	Command
Display neighbor discovery information.	clear ipv6 neighbors on page 635
	clear ipv6 router-advertisement on page 636
	show ipv6 neighbors on page 636
	show ipv6 router-advertisement on page 637

clear ipv6 neighbors

Syntax clear ipv6 neighbors (all | host *host-name*)

Description Clear IPv6 router advertisement counters.

Options all—Clear IPv6 neighbors.

host *host-name*—Clear the specificed IPv6 neighbors.

Required Privilege Level view

See Also show ipv6 neighbors on page 636

clear ipv6 router-advertisement

Syntax clear ipv6 router-advertisement <interface *interface*>

Description Clear IPv6 router advertisement counters.

Options none—Clear IPv6 router advertisement counters for all interfaces.

interface *interface*—(Optional) Clear IPv6 router advertisement counters for the specified interface.

Required Privilege Level view

See Also show ipv6 router-advertisement on page 637

show ipv6 neighbors

Syntax show ipv6 neighbors

Description Display information about the IPv6 neighbor cache.

Required Privilege Level view

Output Fields IPv6 Address—Name of IPv6 interface.

Linklayer Address—Link-layer address.

State—State of link. It can be up, down, reachable, or unreachable.

Exp—Number of seconds until entry expires.

Rtr—Neighbor is a router. It can be yes or no.

Interface—Name of interface.

Sample Output

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp  Rtr  Interface
fe80::2a0:c9ff:fe5b:4c1e     00:a0:c9:5b:4c:1e  reachable   15   yes  fxp0.0
```

show ipv6 router-advertisement

Syntax show ipv6 router-advertisement <conflicts> <interface *interface*><prefix *prefix*/*prefix length*>

Description Display information about about IPv6 router advertisements. The output contains statistics about messages sent and received on interfaces, as well as information received from advertisements from other routers.

Options conflicts—Display only the IPv6 router advertisement information that is conflicting.

interface *interface*—Display IPv6 router advertisement information for the specified interface.

prefix *prefix*/*prefix length*—Display IPv6 router advertisement information for the specified prefix.

Required Privilege Level view

See Also clear ipv6 router-advertisement on page 636

Output Fields Interface—Name of interface.

Advertisements sent—Number of router advertisements sent and time since they were sent.

Solicits received—Number of solicitation messages received.

Advertisements received—Number of router advertisements received.

Advertisements from—Names of interfaces from which router advertisements have been received and time since the last one was received.

Managed—Managed address configuration flag. It can be 0 (stateless) or 1 (stateful).

Other configuration—Other stateful configuration flag. It can be 0 (stateless) or 1 (stateful).

Reachable time—The time that a node assumes a neighbor is reachable after receiving a reachability confirmation (in milliseconds).

Default lifetime—Default lifetime (in seconds). It can be from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.

Retransmit timer—The time between retransmitted Neighbor Solicitation messages (in milliseconds).

Current hop limit—Configured current hop limit. It can be 1 to 255.

Prefix—Name and length of prefix.

Valid lifetime—How long the prefix remains valid for onlink determination. By default, the valid lifetime is set to 2,592,000 seconds. When 0xffffffff is specified, the valid lifetime is infinite. The valid lifetime is always greater than the preferred lifetime.

Preferred lifetime—How long the prefix generated by stateless autoconfiguration remains preferred. By default, the preferred lifetime is set to 604,800 seconds. When preferred lifetime is set to 0xffffffff, the preferred lifetime is infinite. The preferred lifetime is never greater than the valid lifetime.

On link—Onlink flag. It can be 0 (not onlink) or 1 (onlink). The default is 1.

Autonomous—Autonomous address configuration flag. It can be 0 (not autonomous) or 1 (autonomous). The default is 1. Conflicting information is denoted by displaying the value the router is advertising in brackets.

Sample Output: show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

Sample Output: show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

Sample Output: show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

Sample Output: show ipv6 router-advertisement no conflicts

```
user@host> show ipv6 router-advertisement conflicts
```

Chapter 41

IP Multicast Monitoring and Troubleshooting

Table 81 summarizes the command-line interface (CLI) commands you can use to monitor IP multicast. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 81: Commands for Monitoring IP Multicast

Task Category	Task or Information to Monitor	Command
Generic IP Multicast Information	Entries in the multicast forwarding cache	show multicast route on page 666
	Entries in the multicast next-hop table	show multicast nexthops on page 665
	Multicast reverse-path-forwarding calculations	show multicast rpf on page 669
	Administratively scoped addresses	show multicast scope on page 671
	Announced multicast sessions	show multicast sessions on page 672
	Multicast statistics	show multicast statistics on page 673
	Clear multicast statistics	clear multicast statistics on page 644
	Most active multicast groups	show multicast usage on page 675
	Trace information about a multicast path	mtrace on page 647 mtrace from-source on page 648 mtrace monitor on page 650 mtrace to-gateway on page 651
	Multicast network configuration information.	show multicast mrinfo on page 664
DVMRP	DVMRP and PIM tunnels	show multicast tunnels on page 674
	DVMRP neighbors	show dvmrp neighbors on page 654
	DVMRP prefixes	show dvmrp prefix on page 655
	DVMRP prunes	show dvmrp prunes on page 657
	Status of interfaces on which DVMRP is configured	show dvmrp interfaces on page 653
	Most active DVMRP groups	show multicast usage on page 675
MSDP	MSDP peers	show msdp on page 662
	Clear MSDP source active cache	clear msdp cache on page 644
	MSDP statistics	show msdp statistics on page 663
	Clear MSDP statistics	clear msdp statistics on page 644
	Test MSDP peers	test msdp on page 688

Task Category	Task or Information to Monitor	Command
PIM	Bootstrap routers	show pim bootstrap on page 676
	PIM groups	show pim join on page 679
	Clear the PIM join and prune states	clear pim join on page 645
	Status of interfaces on which PIM is configured	show pim interfaces on page 677
	PIM neighbors	show pim neighbors on page 682
	Rendezvous points	show pim rps on page 683
	PIM source RPF state	show pim source on page 686
	PIM (*,*,RP) join and prune states	show sap listen on page 688
	PIM statistics	show pim statistics on page 687
	Clear PIM statistics	clear pim statistics on page 646
	Most active PIM groups	show multicast usage on page 675
IGMP	Members of IGMP groups	show igmp group on page 658
	Members of IGMP groups by interface	show igmp interface on page 660
	Clear IGMP group members	clear igmp membership on page 641
	IGMP statistics	show igmp statistics on page 661
	Clear IGMP statistics	clear igmp statistics on page 643
SAP / SDP	Session announcement addresses	show sap listen on page 688

clear igmp membership

Syntax clear igmp membership <interface *interface-name*> <group *address-range*>

Description Clear IGMP group members.

Options none—Clear all IGMP members on all interfaces and for all address ranges.

group *address-range*—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is 224.2/16. If you omit the destination prefix length, the default is /32.

interface *interface-name*—(Optional) Clear all IGMP group members on an interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

Required Privilege Level clear

See Also show igmp group on page 658, show igmp interface on page 660

Sample Output

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          186
so-0/0/0         224.2.127.254   10.1.128.1          186
so-0/0/0         239.255.255.255 10.1.128.1          187
so-0/0/0         224.1.127.255   10.1.128.1          188
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

After some time, the groups are relearned:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

```
user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

After some time, the groups are relearned:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.1.127.255   10.1.128.1          231
so-0/0/0         224.2.127.254   10.1.128.1          233
so-0/0/0         224.2.127.253   10.1.128.1          236
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership group 239.225/16 interface so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.1.127.255   10.1.128.1          231
so-0/0/0         224.2.127.254   10.1.128.1          233
so-0/0/0         224.2.127.253   10.1.128.1          236
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp statistics

Syntax clear igmp statistics <interface *interface-name*>

Description Clear IGMP statistics.

Options none—Clear IGMP statistics on all interfaces.

interface *interface-name*—(Optional) Clear IGMP statistics on a specific interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

Required Privilege Level clear

See Also show igmp statistics on page 661

Sample Output

```
user@host> show igmp statistics
IGMP statistics on all interfaces:
IGMP Message type     Received       Sent  Rx errors
Queries                    997      18296          0
Report V1                    0          0          0
DVMRP                    19784      35476          0
PIM V1                   18310          0          0
Cisco Trace                  0          0          0
Report V2                    6          0          0
Leave                        0          0          0

Domain Wide Reports          0          0          0

IGMP statistics summary for all interfaces:
Unknown type                 0
Bad Length                   0
Bad Checksum                 0
Bad Receive If               0
Rx non-local                 0

user@host> clear igmp statistics

user@host> show igmp statistics
IGMP statistics on all interfaces:
IGMP Message type     Received       Sent  Rx errors
Queries                      0          2          0
Report V1                    0          0          0
DVMRP                        2          3          0
PIM V1                       0          0          0
Cisco Trace                  0          0          0
Report V2                    0          0          0
Leave                        0          0          0
Mtrace Response              0          0          0
Mtrace Request               0          0          0
Domain Wide Reports          0          0          0

IGMP statistics summary for all interfaces:
Unknown type                 0
Bad Length                   0
Bad Checksum                 0
Bad Receive If               0
Rx non-local                 0
```

clear msdp cache

Syntax clear msdp cache <peer *peer address*>

Description Clear the entries in the MSDP source active cache.

Options none—Clear all entries in the MSDP source active cache.

peer *peer address*—(Optional) Clear the MSDP source active cache entries learned from a specific peer. You can include wildcard characters in the peer name, as described in "Wildcard Characters in Interface Names" on page 187.

Required Privilege Level clear

See Also show msdp on page 662

clear msdp statistics

Syntax clear msdp statistics <peer *peer address*>

Description Clear MSDP statistics.

Options none—Clear all MSDP statistics.

peer *peer address*—(Optional) Clear the MSDP statistics learned from a specific peer. You can include wildcard characters in the peer name, as described in "Wildcard Characters in Interface Names" on page 187.

Required Privilege Level clear

See Also show msdp statistics on page 663

clear multicast statistics

Syntax clear multicast statistics <interface *interface-name*>

Description Clear multicast statistics.

Options none—Clear multicast statistics on all interfaces.

interface *interface-name*—(Optional) Clear multicast statistics on a specific interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

Required Privilege Level clear

See Also show multicast statistics on page 673

clear pim join

Syntax clear pim join *<group-address>*

Description Clear the PIM join and prune states.

Options none—Clear the PIM join and prune states for all groups.

group-address—(Optional) Clear a group address.

Required Privilege Level clear

See Also show pim join on page 679

clear pim statistics

Syntax clear pim statistics <interface *interface-name*>

Description Clear PIM statistics on the specified interface.

Options interface *interface-name*—Name of an interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 110.

Required Privilege Level clear

See Also show pim statistics on page 687

Sample Output

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                   2111       4222          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune             14200      13115          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0

PIM statistics summary for all interfaces:
Unknown type                            0
V1 Unknown type                         0
Unknown Version                         0
Neighbor unknown                        0
Bad Length                              0
Bad Checksum                            0
Bad Receive If                          0
Rx Intf disabled                     2007
Rx V1 Require V2                        0
Rx Register not RP                      0
RP Filtered Source                      0
Unknown Reg Stop                        0
Rx Join/Prune no state               1040
Rx Graft/Graft Ack no state             0
...
```

```
user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                      1          0          0
V1 Register                   0          0          0
...
```

mtrace

Syntax mtrace <*source*>

Description Display trace information about a multicast path. The mtrace command for multicast traffic is similar to the traceroute command used for unicast traffic. Unlike traceroute, mtrace traces traffic backwards, from the receiver to the source.

Options *source*—Source host name or address.

Required Privilege Level view

Output Fields hop—Number of hops from source to the listed router.

router name—Name of router for this hop.

address—Address of router for this hop.

protocol—Protocol used.

ttl—TTL threshold.

Sample Output

```
user@host> mtrace 192.1.4.2
mtrace: WARNING: no multicast group specified, so no statistics printed
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

mtrace from-source

Syntax mtrace from-source source < *source*> <extra-hops *extra-hops* | group *group* | interval *interval* | max-hops *max-hops* | max-queries *max-queries* | response *response* | ttl *ttl* | wait-time *wait-time*> <loop> <multicast-response | unicast-response> <no-resolve> <no-router-alert> <brief | detail>

Description Display trace information about a multicast path from a source to this router. If a group address is given, additional information, such as packet rates and losses, can be gathered.

Options brief—(Optional) Do not display packet rates and losses.

detail—(Optional) Display packet rates and losses if group address is given.

extra-hops *extra-hops*—(Optional) Number of hops after reaching a non-responding router. It can be 0 through 255.

group *group*—(Optional) Group address for which to trace the path. The default group address is 0.0.0.0.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default is 10 seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward source. It can be 0 through 255. The default is 32 hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. It can be 1 through 32. The default is 3.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to host or multicast address.

source *source*—Source host name or address.

ttl *ttl*—(Optional) IP time-to-live value. It can be 0 through 255. Local queries to the multicast group use TTL 1, otherwise the default is 127.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default is 3 seconds.

Required Privilege Level view

Output Fields

hop—Number of hops from source to the listed router.

router name—Name of router for this hop.

address—Address of router for this hop.

protocol—Protocol used.

ttl—TTL threshold.

source—Source address.

Response Dest—Response destination address.

Overall—Average packet rate for all traffic at each hop.

Packet Statistics for Traffic From—Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.

Receiver—IP address receiving the multicast.

Query Source—IP address sending the mtrace query.

Sample Output

```
user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.


Waiting to accumulate statistics...Results after 10 seconds:


  Source        Response Dest    Overall     Packet Statistics For Traffic From
192.1.4.2 192.1.1.2    Packet      192.1.4.2 To 225.1.1.1
     v       __/  rtt    2 ms     Rate       Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2  routerC.lab.mycompany.net
     v     ^      ttl    2                      0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
     v      \__   ttl    3                      ?/0            0 pps
192.1.1.2   192.1.1.2
  Receiver      Query Source
```

mtrace monitor

Syntax mtrace monitor

Description Listen passively for multicast responses.

To exit mtrace monitor, type Ctrl+C.

Required Privilege Level view

Output Fields Mtrace query at—Date and time of query.

by—Address of host issuing query.

resp to—Response destination.

qid—Query ID number.

packet from...to—IP address of query source and default group destination.

from...to—IP address of multicast source and response address.

via group—IP address of group to trace.

mxhop—Maximum hop setting. It can be 0 through 255. The default is 32.

Sample Output

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)


Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)


Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)


Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

mtrace to-gateway

Syntax mtrace to-gateway gateway < *gateway*> <brief | detail> <extra-hops *extra-hops* | group *group* | interface *interface* | interval *interval* | max-hops *max-hops* | max-queries *max-queries* | response *response* | ttl *ttl* | wait-time *wait-time*> <loop> <multicast-response | unicast-response> <no-resolve> <no-router-alert>

Description Display trace information about a multicast path from this router to gateway router.

Options brief—(Optional) Do not display packet rates and losses.

detail—(Optional) Display packet rates and losses if group address is given.

extra-hops *extra-hops*—(Optional) Number of hops after reaching a nonresponding router. It can be 0 through 255.

gateway *gateway*—Send trace query to gateway multicast address.

group *group*—(Optional) Group address for which to trace the path. The default group address is 0.0.0.0.

interface *interface*—(Optional) Source address for sending trace query.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. Default is 10 seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward source. It can be 0 through 255. The default is 32 hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. It can be 1 through 32. The default is 3.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to host or multicast address.

ttl *ttl*—(Optional) IIP time-to-live value. It can be 0 through 255. Local queries to the multicast group use TTL 1, otherwise the default is 127.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default is 3 seconds.

Required Privilege Level view

Output Fields hop—Number of hops from source to the listed router.

router name—Name of router for this hop.

address—Address of router for this hop.

protocol—Protocol used.

ttl—TTL threshold.

Packet Statistics for Traffic From—Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.

Sample Output

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1
interface 192.1.1.73 brief
Mtrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerA.lab.mycompany.net (192.1.1.2)  PIM  thresh^ 1
 -2  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -3  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

show dvmrp interfaces

Syntax show dvmrp interfaces <brief>

Description Display information about the interfaces on which DVMRP is enabled.

Options brief—(Optional) Display brief information about the DVMRP interfaces.

Default: brief

Required Privilege Level view

Output Fields Interface—Name of the interface.

State—State of the interface. It can be Up or Down.

Leaf—Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors.

Metric—Interface metric. It can be a value from 1 through 31.

Announce—Number of routes the interface is announcing.

Mode—DVMRP mode. It can be Forwarding or Unicast-routing. Default is Forwarding.

- Forwarding—DVMRP does both the routing and the multicast data forwarding.
- Unicast-routing—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface.

Sample Output

```
user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0     Up     N     1     4 Forwarding
fxp1.0     Up     N     1     4 Forwarding
fxp2.0     Up     N     1     3 Forwarding
lo0.0      Up     Y     1     0 Unicast-routing
```

show dvmrp neighbors

Syntax show dvmrp neighbors <brief >

Description Display information about DVMRP neighbors.

Options brief—(Optional) Display brief information about DVMRP neighbors.

Default: brief

Required Privilege Level view

Output Fields Neighbor—Address of the neighboring DVMRP router.

Interface—Interface through which the neighbor is reachable.

Version—Version of DVMRP that the neighbor is running, in the format *major.minor*.

Flags—Information about the neighbor:

- 1—One way. The local router has seen the neighbor, but the neighbor has not seen the local router.
- G—Neighbor supports generation ID.
- L—Neighbor is a leaf router.
- M—Neighbor supports mtrace.
- N—Neighbor supports netmask in prunes and grafts.
- P—Neighbor supports pruning.
- S—Neighbor supports SNMP.

Routes—Number of routes learned from the neighbor.

Timeout—How long until the DVMRP neighbor information times out, in seconds.

Transitions—Number of generation ID changes that have occurred since the local router learned about the neighbor.

Sample Output: show dvmrp neighbors

```
user@host> show dvmrp neighbors
Neighbor        Interface        Version Flags   Routes Timeout Transitions
192.168.1.1    ipip.0             3.255  PGM         3      28           1
```

Sample Output: show dvmrp neighbors (PIM not running)

```
user@host> show dvmrp neighbors
PIM is not running
```

show dvmrp prefix

Syntax show dvmrp prefix <*prefix*> <brief >

Description Display information about DVMRP prefixes.

Options brief—(Optional) Display brief information about DVMRP prefixes.

prefix—(Optional) Display information about specific prefixes.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show dvmrp prefix brief on page 656
Sample Output: show dvmrp prefix detail on page 656

Options at a Glance Table 82 summarizes which information is included in each show dvmrp prefix command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 82: Show DVMRP Prefix Output Field Summary

Options	Field Description
All	Age—Last time that the route was refreshed.
Detail	Cache lifetime—Lifetime of the group in the multicast cache, in seconds.
Detail	Grafts sent—Number of grafts sent to the multicast group.
Detail	*group*—Groups seen by the prefix.
All	Next hop—Next hop from which the route was learned.
All	Prefix—DVMRP route.
Detail	Prunes lifetime—Lifetime remaining and total lifetime of prunes, in seconds.
Detail	Prunes sent—Number of prunes sent to the multicast group.

Output Fields Prefix—DVMRP route.

Next hop—Next hop from which the route was learned.

Age—Last time that the route was refreshed.

group—(Detail output only) Groups seen by the prefix.

Prunes sent—(Detail output only) Number of prunes sent to the multicast group.

Grafts sent—(Detail output only) Number of grafts sent to the multicast group.

Cache lifetime—(Detail output only) Lifetime of the group in the multicast cache, in seconds.

Prune lifetime—(Detail output only) Lifetime remaining and total lifetime of prunes, in seconds.

Sample Output: show dvmrp prefix brief

```
user@host> show dvmrp prefix brief
Prefix             Next hop        Age
0.0.0.0        /0  192.168.1.1     00:00:34
127.0.0.1      /32 127.0.0.1       16:26:50
128.109.46.0   /24 192.168.1.1     00:00:34
192.168.1.0    /24 192.168.1.1     00:00:34
204.70.248.160 /29 204.70.248.161  16:26:50
204.70.248.168 /29 204.70.248.162  16:26:49
204.70.248.176 /28 204.70.248.187  16:26:50
204.70.248.192 /26 204.70.248.224  16:26:50
```

Sample Output: show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix             Next hop        Age
0.0.0.0        /0  192.168.1.1     00:00:03
    224.2.127.254:
      Prunes sent Grafts sent Cache lifetime Prune lifetime
               0          0          0          0/    0
    224.2.127.254:
      Prunes sent Grafts sent Cache lifetime Prune lifetime
               0          0          0          0/    0
    224.2.139.68:
      Prunes sent Grafts sent Cache lifetime Prune lifetime
               2          2          0          0/ 7200
    224.2.245.3:
      Prunes sent Grafts sent Cache lifetime Prune lifetime
               2          2          0          0/ 7200
...
```

Sample Output: show dvmrp prefix (PIM not running)

```
user@host> show dvmrp prefix
PIM is not running
```

show dvmrp prunes

Syntax show dvmrp prunes <all> <rx> <tx> <brief>

Description Display information about active DVMRP prunes.

Options all—(Optional) Display information about all received and transmitted prunes.

brief—(Optional) Display brief information about DVMRP pruning.

rx—(Optional) Display information about received prunes.

tx—(Optional) Display information about transmitted prunes.

Default: brief

Required Privilege Level view

Output Fields Group—Group address.

Source prefix—Prefix that the prune is for.

Timeout—How long until the prune message expires, in seconds.

Neighbor—Neighbor to whom the prune was sent or from whom the prune was received.

Sample Output

```
user@host> show dvmrp prunes
Group           Source prefix      Timeout Neighbor
224.0.1.1       128.112.0.0    /12    7077 192.168.1.1
224.0.1.32      160.0.0.0      /3     7087 192.168.1.1
224.2.123.4     136.0.0.0      /5     6955 192.168.1.1
224.2.127.1     129.0.0.0      /8     7046 192.168.1.1
224.2.135.86    128.102.128.0  /17    7071 192.168.1.1
224.2.135.86    129.0.0.0      /8     7074 192.168.1.1
224.2.135.86    130.0.0.0      /7     7071 192.168.1.1
...
```

Sample Output: show dvmrp prunes (PIM not running)

```
user@host> show dvmrp prunes
PIM is not running
```

show igmp group

Syntax show igmp group <*group-name*> <brief | detail>

Description Display information about IGMP group membership.

Options brief—(Optional) Display brief information about IGMP group members.

detail—(Optional) Display detailed information about IGMP group members.

group-name—(Optional) IP address of an IGMP group.

Default: brief

Required Privilege Level view

See Also clear igmp membership on page 641

Output Fields Interface—Name of the interface which received the IGMP membership report. A name of "local" means that the local router joined the group itself.

Group—Group address.

Source—Source address.

Last reported—Address of the host that joined the group.

Timeout—When the group membership expires.

Sample Output

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
ge-4/1/0.0       224.1.1.1       10.6.0.7            239
ge-4/1/0.0       224.1.1.2       10.6.0.7            252
ge-4/1/0.0       224.1.1.3       10.6.0.7            257
ge-5/3/0.0       224.3.3.1       10.9.0.5            138
ge-5/3/0.0       224.3.3.2       10.9.0.5            227
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
local            224.0.1.39      (null)                0
local            224.0.1.40      (null)                0
```

The following example shows output from an IGMPv3 enabled router.

```
user@host> show igmp group
Interface Group Source Last Reported Timeout
fxp0.0 224.0.0.2 0.0.0.0 192.168.14.144 193
fxp0.0 224.0.0.22 0.0.0.0 192.168.14.142 193
fe-2/0/1.0 232.232.10.10 192.168.195.42 192.168.195.34 0
fe-2/0/1.0 232.232.10.10 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.11 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.12 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.13 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.14 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.15 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.16 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.17 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.18 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.19 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.20 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.21 192.168.195.62 192.168.195.34 0
fe-2/0/1.0 232.232.10.30 192.168.195.94 192.168.195.34 0
fe-2/0/1.0 232.232.10.31 192.168.195.94 192.168.195.34 0
fe-2/0/1.0 232.232.10.32 192.168.195.94 192.168.195.34 0
fe-2/0/1.0 232.232.10.33 192.168.195.94 192.168.195.34 0
fe-2/0/1.0 232.232.10.34 192.168.195.94 192.168.195.34 0
local 224.0.0.2 0.0.0.0 0.0.0.0 0
```

show igmp interface

Syntax show igmp interface <*interface-name*> <brief | detail>

Description Display information about the interfaces on which IGMP is configured.

Options brief—(Optional) Display brief information about IGMP interfaces.

detail—(Optional) Display detailed information about IGMP interfaces.

interface-name—(Optional) Display information about the specified interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 110.

Default: brief

Required Privilege Level view

See Also clear igmp membership on page 641

Output Fields Interface—Name of the interface.

State—State of the interface. It can be Up or Down.

Querier—Address of the neighbor that has been elected to send membership queries.

Timeout—How long until the IGMP querier is declared to be unreachable, in seconds.

Version—IGMP version being used on the interface. It can be 1, 2, or 3.

Groups—Number of groups on the interface.

Sample Output

```
user@host> show igmp interface
Interface        State Querier         Timeout Version Groups
ge-2/0/0.0       Down                        0       2      0
ge-4/0/0.0       Up    10.3.0.9            199       2      0
ge-4/1/0.0       Up    10.6.0.1           None       2      3
ge-3/0/0.0       Down                        0       2      0
ge-5/3/0.0       Up    10.9.0.1           None       2      2
```

show igmp statistics

Syntax show igmp statistics <interface *interface-name*>

Description Display IGMP statistics.

Options interface *interface-name*—(Optional) Display statistics about the specified interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 110.

Required Privilege Level view

See Also clear igmp statistics on page 643

Output Fields IGMP statistics on all interfaces or IGMP statics for interfaces *interface-name*—Name of the interface for which the statistics are being reported.

Received—Number of received statistics.

Transmitted—Number of transmitted statistics.

Rx errors—Number of received packets that contained errors.

IGMP statistics summary for all interfaces—Summary of IGMP statistics for all interfaces.

Sample Output

```
user@host> show igmp statistics
IGMP statistics on all interfaces:
IGMP message type     Received Transmitted  Rx errors
Queries                    986      18131          0
Report V1                    0          0          0
DVMRP                    19603      35171          0
PIM V1                   18137          0          0
Cisco Trace                  0          0          0
Report V2                    6          0          0
Leave                        0          0          0
Mtrace Response              0          0          0
Mtrace Request               0          0          0
Domain Wide Reports          0          0          0

IGMP statistics summary for all interfaces:
Unknown type                 0
Bad Length                   0
Bad Checksum                 0
Bad Receive If               0
Rx non-local                 0
```

show msdp

Syntax show msdp <peer *peer address*> <source-active *group* | *originator* | *source*> <brief | detail>

Description Display information about MSDP.

Options brief—(Optional) Display brief information about MSDP.

detail—(Optional) Display detailed information about MSDP.

peer *peer address*—(Optional) Display information about a particular MSDP peer.

source-active *group*—(Optional) Display source active cache information for a particular group.

source-active *originator*—(Optional) Display information about the MSDP peer that originated the source active cache entries.

source-active *source*—(Optional) Display source active cache for a particular source.

Required Privilege Level view

Output Fields Peer—Address of peer.

Local address—Local address of peer.

State—Status of the MSDP connection. It can be Listen, Established, or Inactive.

Last up/down—Time of most recent peer-state change.

Peer Connect Retries—Number of peer connect retries.

State timer expires—Number of seconds before state timer expires.

Peer Times out—Number of peer timeouts.

Sample Output: show msdp (standard)

```
user@host> show msdp
Peer address Local address State Last up/down Peer-Group
10.255.245.34 10.255.245.39 Listen 00:05:07
```

Sample Output: show msdp detail

```
user@host> show msdp detail
Peer: 10.255.245.34
Local address: 10.255.245.39
State: Listen
Peer Connect Retries: 0
State timer expires: 0
Peer: 10.255.245.43
Local address: 10.255.245.39
State: Established
Peer Connect Retries: 0
State timer expires: 53
Peer Times out: 2
```

show msdp statistics

Syntax show msdp statistics <peer *peer address*>

Description Display statistics about MSDP peers.

Options peer *peer address*—(Optional) Display information about a particular MSDP peer.

Output Fields Peer—Address of peer.

Last State Change—How long ago the peer state changed.

Last message received from peer—How long ago the last message was received from peer.

RPF Failures—Number of RPF failures.

Remote Closes—Number of times the remote peer closed.

Peer Timeouts—Number of peer timeouts.

SA messages sent—Number of source active messages sent.

SA messages received—Number of source active messages received.

SA request messages sent—Number of source active request messages sent.

SA request messages received—Number of source active request messages received.

SA response messages sent—Number of source active response messages sent.

SA response messages received—Number of source active response messages received.

Keepalive messages sent—Number of keepalive messages sent.

Keepalive messages received—Number of keepalive messages received.

Unknown messages received—Number of unknown messages received.

Error messages received—Number of unknown messages received.

Sample Output

```
user@host> show msdp statistics
Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0
```

show multicast mrinfo

Syntax show multicast mrinfo <*address* | *dns-device-name*>

Description Display configuration information about multicast networks, including neighboring multicast router addresses.

Options none—Display configuration information about all multicast networks.

address—(Optional) Display configuration information about a particular neighbor.

dns-device-name—(Optional) Display configuration information about a particular DNS host.

Required Privilege Level view

Output Fields Source address information—Query address, host name (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.

IP address 1—Queried router interface address.

IP address 2—Directly attached neighbor interface address.

Name—Name or IP address of neighbor.

[metric/threshold/type/flags]—Neighbor's multicast profile:

- metric—Always has a value of 1, because mrinfo queries the directly connected interfaces of a device.
- threshold—Multicast threshold (ttl). It can have a value between 0 and 255.
- type—Multicast connection type. It can be pim or tunnel.
- flags—Flags for this route.
  - querier—Indicates the queried router is the designated router for the neighboring session.
  - leaf—Indicates the link is a leaf in the multicast network.
  - down—Link status indicator.

Sample Output

```
user@host>show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast nexthops

Syntax show multicast nexthops <identifier *identifier-number*> <brief | detail>

Description Display the entries in the multicast next-hop table.

Options brief—(Optional) Display brief information about the entries.

detail—(Optional) Display detailed information about the entries.

identifier *identifier-number*—(Optional) Show particular next-hop by ID number. It can be a number between 1 and 65535.

Default: brief

Required Privilege Level view

Output Fields ID—Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine.

Refcnt—Number of cache entries that are using this next hop.

KRefcount—The kernel's reference count for the next hop.

Downstream interface—Interface names associated with each multicast nexthop ID.

Sample Output

```
user@host> show multicast nexthops brief
ID      Refcount  KRefcount Downstream interface
11             2          1 local
                             so-1/1/0.0

33             4          2 so-0/0/0.0
                             so-3/1/0.0
                             t3-4/2/0.0
                             ge-6/3/0.0

39             8          4 so-0/0/0.0
                             at-2/1/0.0
                             at-2/1/0.1
                             so-3/1/0.0
                             t3-4/2/0.0
                             ge-6/3/0.0

62             6          3 so-0/0/0.0
                             at-2/1/0.0
                             at-2/1/0.1
                             so-3/1/0.0
                             ge-6/3/0.0

79            78         39 so-0/0/0.0
                             at-2/1/0.0
                             so-3/1/0.0
                             ge-6/3/0.0
                             ge-7/2/0.0
                             ge-7/3/0.0
```

show multicast route

Syntax show multicast route <*regular-expression*> <all> <active> <inactive> <group *group*> <source-prefix *prefix*> <brief | detail | extensive> < *instance* >

Description Display the entries in the multicast forwarding table. You can display similar information with the show route table inet.1 command.

Options active—(Optional) Display all active entries in the multicast forwarding table.

all—(Optional) Display all entries in the multicast forwarding table.

brief—(Optional) Display brief information about the multicast forwarding table entries.

detail—(Optional) Display detailed information about the multicast forwarding table entries.

group *group*—(Optional) Display the cache entries for a particular group.

inactive—(Optional) Display all inactive entries in the multicast forwarding table.

instance—(Optional) Display all interfaces under the named routing instance.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.

source-prefix *prefix*—(Optional) Display the cache entries for a particular source prefix.

Default: brief

Required Privilege Level view

Options at a Glance Table 83 summarizes which information is included in the output of each show multicast route command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 83: Show Multicast Route Output Field Summary

Options	Field Description
All	Act—Whether the group is active (A) or inactive (I).
Detail	AgeOut—Time until the prefix is removed from the multicast forwarding table.
Detail	Cloned from—Actual prefix on which the multicast forwarding table entry is based.
Extensive	Forwarding rate—Rate at which packets are being forwarded for this source and group entry (in pps, kbps).
All	Group—Group address.
All	InIf—Index of the incoming interface.
All	Instance—Name of the routing instance.
All	NHid—Next-hop identifier of the prefix.
Detail Extensive	Packets—Number of packets that have been forwarded to this prefix.
All	Pru—Whether the prefix is pruned (P) or forwarding (F).
All	Session name—Name of the multicast session.
All	Source prefix—Prefix and length of the source as it is in the multicast forwarding table.
Detail Extensive	Source—Source address of the host that triggered creation of the multicast forwarding table entry.
Extensive	Upstream interface—(Extensive only) Name of the interface on which the packet with this source prefix is expected to arrive.

Output Fields

Instance—Name of the routing instance.

Group—Group address.

Source prefix—Prefix and length of the source as it is in the multicast forwarding table.

Act—Whether the group is active (A) or inactive (I).

Pru—Whether the prefix is pruned (P) or forwarding (F).

InIf—Index of the incoming interface. The index also is listed in the output of the show interfaces command.

NHid—Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command.

AgeOut—Time until the prefix is removed from the multicast forwarding table.

Packets—Number of packets that have been forwarded to this prefix.

Session—Name of the multicast session.

Cloned from—Actual prefix on which the multicast forwarding table entry is based.

Source—Source address of the host that triggered creation of the multicast forwarding table entry.

Session name—Name of the multicast session.

Upstream interface—Name of the interface on which the packet with this source prefix is expected to arrive.

Forwarding rate— Rate at which packets are being forwarded for this source and group entry (in pps, kbps).

Sample Output: show multicast route brief

```
user@host> show multicast route brief
show multicast route
Group           Source prefix     Act Pru InIf  NHid  Session Name
239.1.1.1       10.255.71.46   /32 A   F  0     37    Administratively Scoped
```

Sample Output: show multicast route detail

```
user@host> show multicast route detail instance VPN-A
show multicast route detail instance VPN-A
Group           Source prefix     Act Pru InIf  NHid  Packets    Session
225.1.1.1       192.168.195.38 /32 A   F  14    48    396        MALLOC
```

Sample Output: show multicast route extensive

```
user@host> show multicast route extensive instance VPN-A
show multicast route extensive instance VPN-A
Group           Source prefix     Act Pru NHid  Packets    IfMismatch Timeout
225.1.1.1       192.168.195.38 /32 A   F  48    529        2          360
    Upstream interface: mt-1/1/0.32769
    Session name: MALLOC
    Forwarding rate: 0 kbps (1 pps)
```

show multicast rpf

Syntax show multicast rpf <*prefix*> <brief | detail> <*dns-name*> < *instance* >

Description Display information about multicast reverse-path-forwarding calculations.

Options brief—(Optional) Display brief information about the RPF calculations.

detail—(Optional) Display detailed information about the RPF calculations.

dns-name—(Optional) Display the RPF calculation information for the specified DNS hostname.

prefix—(Optional) Display the RPF calculation information for the specified prefix.

instance—(Optional) Display all interfaces under the named routing instance.

Required Privilege Level view

Options at a Glance Table 84 summarizes which information is included in each of the show multicast rpf command options. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 84: Show Multicast RPF Output Field Summary

Options	Field Description
Detail	Cloned from—Actual prefix on which the multicast forwarding table entry is based.
All	Instance—Name of the routing instance.
All	Protocol—How the route was learned.
All	RPF interface—Upstream interface.
All	RPF neighbor—Upstream neighbor.
All	Source prefix—Prefix and length of the source as it is in the multicast forwarding table.

Output Fields Instance—Name of the routing instance.

Source prefix—Source route.

Protocol—How the route was learned.

RPF interface—Upstream interface.

RPF neighbor—Upstream neighbor.

Cloned from—(Detail output only) Actual prefix on which the cache entry is based.

Sample Output: show multicast rpf (standard, instance)

```
user@host>show multicast rpf instance VPN-A
Multicast RPF table: INET.1

Source prefix      Protocol   RPF interface    RPF neighbor
10.255.14.30/32    OSPF       t1-1/0/0:0.0     (null)
10.255.245.91/32   BGP        so-0/0/2.0
192.168.195.32/30  OSPF       t1-1/0/0:0.0     (null)
192.168.195.36/30  BGP        so-0/0/2.0
192.168.196.104/30 Direct     t1-1/0/0:0.0
192.168.196.105/32 Local
224.0.0.2/32       PIM
224.0.0.5/32       OSPF
224.0.0.13/32      PIM
```

Sample Output: show multicast rpf (standard)

```
user@host> show multicast rpf
show multicast rpf
Multicast RPF table: INET.0

Source prefix      Protocol   RPF interface    RPF neighbor
0.0.0.0/0          Static
10.255.71.46/32    Direct     lo0.0
10.255.71.47/32    OSPF       so-0/0/2.0       (null)
10.255.71.51/32    OSPF       so-0/0/2.0       (null)
127.0.0.1/32
192.16.0.0/12      Static     fxp0.0           192.168.71.254
192.168.0.0/18     Static     fxp0.0           192.168.71.254
192.168.40.0/22    Static     fxp0.0           192.168.71.254
192.168.64.0/18    Static     fxp0.0           192.168.71.254
192.168.71.0/24    Direct     fxp0.0
192.168.71.46/32   Local
192.168.196.32/30  OSPF       so-0/0/2.0       (null)
192.168.196.68/30  Direct     so-0/0/2.0
192.168.196.69/32  Local
207.17.136.192/32  Static     fxp0.0           192.168.71.254
224.0.0.2/32       PIM
224.0.0.5/32       OSPF
224.0.0.13/32      PIM
```

Sample Output: show multicast rpf detail

```
user@host> show multicast rpf detail
Source prefix      Protocol   RPF interface    RPF neighbor
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1
0.0.0.0        /2  Cloned     ipip.0           192.168.1.1
Cloned from:
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1

0.0.0.0        /3  Cloned     ipip.0           192.168.1.1
Cloned from:
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1

127.0.0.1      /32 Direct     lo0.0
128.0.0.0      /10 Cloned     ipip.0           192.168.1.1
Cloned from:
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1

128.0.0.0      /11 Cloned     ipip.0           192.168.1.1
Cloned from:
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1

128.32.0.0     /11 Cloned     ipip.0           192.168.1.1
Cloned from:
0.0.0.0        /0  DVMRP      ipip.0           192.168.1.1
...
```

show multicast scope

Syntax show multicast scope <brief | detail>

Description Display administratively scoped multicast information.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

Default: brief

Required Privilege Level view

Output Fields Scope name—Name of the multicast scope.

Group prefix—Range of multicast groups that are scoped.

Interface—Interface that is the boundary of the administrative scope.

Resolve Rejects—Number of kernel resolve rejects.

Sample Output

```
user@host> show multicast scope
                                                                Resolve
Scope name          Group Prefix       Interface                Rejects
“local”             239.255.0.0    /16 ipip.0                  0
```

show multicast sessions

Syntax show multicast sessions <*regular-expression*> <brief | detail |extensive>

Description Display information about announced multicast sessions.

Options brief—(Optional) List the sessions.

detail—(Optional) Format the list of sessions.

extensive—(Optional) Display the list of sessions in a raw format.

regular-expression—(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Default: brief

Required Privilege Level view

Output Fields *session-names*—Name of the known announced multicast sessions.

Sample Output

```
user@host> show multicast sessions
  Bird test
  Colorado/UCSD Networks Class
  Colorado/UCSD wb
  LBL - Deb test
  LabWeb - The Spectro-Microscopy Collaboratory
  NASA - Shuttle Mission STS-95
  NASA TV - Broadcast from NASA HQ
  NLANR Techs Conf. MPEG1 (IP/TV)
  SPT Meeting(private)
  Titanic Trailor
  UC Berkeley MIG seminar
  UW CS&E Colloquium
  test1@fsu-acns
```

show multicast statistics

Syntax show multicast statistics

Description Display multicast statistics.

Required Privilege Level view

See Also clear multicast statistics on page 644

Output Fields Interface—Name of the interface for which statistics are being reported.

Protocol—Primary multicast protocol on the interface.

Resolve primary—Number of resolve requests processed by the primary multicast protocol on the interface.

Resolve secondary—Number of resolve requests processed on the interface by the other protocols configured on the router.

Resolve no route—Number of resolve requests that were ignored because there was no route to the source.

Resolve errors—Number of resolve requests that were ignored because of internal errors.

Interface mismatches—Number of multicast packets that did not arrive on the correct upstream interface.

IfMismatch no route—Number of mismatches that were ignored because there was no route to the source.

IfMismatch errors—Number of mismatches that were ignored because of internal errors.

Sample Output

```
user@host> show multicast statistics
Interface        Protocol      Resolve    Resolve    Resolve    Resolve
                               primary  secondary   no route     errors
local                                0          0          0          0
fxp0.0           PIM-dense        100        100          0          0
fxp1.0           PIM-dense         62         62          0          0
ipip.0           DVMRP           6309       6309          3          0

Interface        Protocol    Interface IfMismatch IfMismatch
                            mismatches   no route     errors
local                                0          0          0
fxp0.0           PIM-dense           0          0          0
fxp1.0           PIM-dense           0          0          0
ipip.0           DVMRP               0          0          0

Resolve requests on interfaces not enabled for multicast 0
Interface Mismatches on interfaces not enabled for multicast 0
```

show multicast tunnels

Syntax show multicast tunnels

Description Display information about DVMRP or PIM tunnels.

Required Privilege Level view

Output Fields Interface—Name of the interface.

Local address—Address of the local side of the tunnel.

Remote address—Address of the remote side of the tunnel.

TTL—Tunnel's time-to-live value. It can be a value from 1 through 255.

Sample Output

```
user@host> show multicast tunnels
Interface        Local address   Remote address  TTL
ipip.0           204.69.248.187  192.168.1.1     64
```

show multicast usage

Syntax show multicast usage <brief | detail><*instance*>

Description Display usage information about the ten most active DVMRP or PIM groups.

Options brief—(Optional) Display brief information about the most active DVMRP or PIM groups.

detail—(Optional) Display detailed information about the most active DVMRP or PIM groups.

instance—(Optional) Display all interfaces under the named routing instance.

Required Privilege Level view

Output Fields Instance—Name of the routing instance.

Group—Group address.

Sources—Number of sources.

Packets—Number of packets that have been forwarded to this prefix.

Bytes—Amount of memory used.

Prefix—IP address.

/len—Prefix length.

Sample Output: show multicast usage (standard, instance)

```
user@host> show multicast usage instance VPN-A
show multicast usage instance VPN-A
Group           Sources Packets              Bytes
224.2.127.254   1       5538                 509496
224.0.1.39      1       13                   624
224.0.1.40      1       13                   624


Prefix          /len Groups Packets              Bytes
192.168.195.34  /32  1      5538                 509496
10.255.14.30    /32  1      13                   624
10.255.245.91   /32  1      13                   624


Group Sources Packets Bytes
224.225.1.1 1 1 0
224.225.1.2 1 1 0
224.225.1.3 1 1 0
224.225.1.4 1 1 0
224.225.1.5 1 1 0
Prefix /len Groups Packets Bytes
10.6.0.1 /32 489 489 0
10.255.245.35 /32 2 0 0
```

show pim bootstrap

Syntax show pim bootstrap <*instance*>

Description For PIM sparse mode only, display information about PIM bootstrap routers.

Options *instance*—(Optional) Display all interfaces under the named routing instance.

Required Privilege Level view

Output Fields Instance—Name of the routing instance.

BSR—Bootstrap router.

Pri—Priority of the router to be elected to be the bootstrap router.

Local address—The local router's address.

Pri—The local router's address priority to be elected as the bootstrap router.

State—The local router's election state. It can be Candidate, Elected, or Ineligible.

Timeout—How long until the local router declares the bootstrap router to be unreachable, in seconds.

Sample Output

```
user@host> show pim bootstrap
Instance: PIM.master

BSR             Pri Local address   Pri State      Timeout
None              0 10.255.71.46      0 InEligible       0
```

Sample Output (standard, instance)

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR             Pri Local address   Pri State      Timeout
None              0 192.168.196.105   0 InEligible       0
```

show pim interfaces

Syntax show pim interfaces <*instance*>

Description Display information about the interfaces on which PIM is configured.

Options *instance*—(Optional) Display all interfaces under the named routing instance.

Required Privilege Level view

Output Fields Instance—Name of the routing instance.

Name—Interface name.

State—State of the interface. The state also is displayed in the show interfaces command.

Mode—PIM mode running on the interface. It can be one of the following:

- Sparse—In sparse mode, routers must join and leave multicast groups explicitly. Upstream routers do not forward multicast traffic to this router unless this router has sent an explicit request (using a join message) to receive multicast traffic.
- Dense—Unlike sparse mode, where data is forwarded only to routers sending an explicit request, dense mode implements a flood-and prune mechanism, similar to DVMRP (the first multicast protocol used to support the MBone).
- Sparse-Dense—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to an RP. Instead, data packets destined for that group are forwarded using PIM-DM rules. A group specified as sparse is mapped to an RP, and data packets are forwarded using PIM-SM rules.

V—PIM version running on the interface. It can be either 1 or 2.

State—State of PIM on the interface. It can be one of the following:

- DR—Designated router.
- NotDR—Not the designated router.
- P2P—Point to point.

Priority—Interface's priority to become the designated router.

DR address—Address of the designated router.

Neighbors—Number of neighbors that have been seen on the interface.

Sample Output

```
user@host> show pim interfaces
Name             Stat Mode        V State   Priority DR address       Neighbors
lo0.0            Up   SparseDense 2 DR             1 10.10.20.101             0
at-1/0/0.0       Up   SparseDense 1 P2P                                       1
so-1/1/0.0       Up   SparseDense 1 P2P                                       1
ge-2/0/0.0       Up   SparseDense 2 DR             3 10.10.20.1               3
ge-3/0/0.0       Up   SparseDense 2 NotDR          1 10.10.30.11              2
so-5/0/0.0       Down SparseDense 1 P2P                                       0
so-5/0/1.0       Up   SparseDense 1 P2P                                       1
so-5/0/2.0       Down SparseDense 1 P2P                                       0
pe-6/3/0.32769   Up   Sparse      2 P2P                                       0
```

show pim join

Syntax show pim join <*range*> <brief | detail | extensive> < *instance* >

Description Display information about PIM groups.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

extensive—(Optional) Display very detailed information.

instance—(Optional) Display all interfaces under the named routing instance.

range—(Optional) Address range of the group, specified as *prefix/prefix-length.*

Default: brief

Required Privilege Level view

See Also clear pim join on page 645

Sample Output Sample Output: show pim join brief on page 681
Sample Output: show pim join detail on page 681
Sample Output: show pim join extensive on page 681

Options at a Glance Table 85 summarizes which information is included in each show pim join command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 85: Show PIM Join Output Field Summary

Options	Field Description
Detail Extensive	Downstream interfaces—Detailed information about downstream interfaces, including interface name, state, flags, and timeout.
Extensive	Downstream neighbors—Detailed information about downstream interfaces, including interface name, state, flags, and timeout.
All	Flags—PIM flags.
All	Group—Group address.
All	Instance—Name of the routing instance.
All	RP—Rendezvous point for the PIM group.
All	Source—Multicast source.
All	Upstream interface—RPF Interface toward the source address for the source specific state (S, G) or toward the RP address for the non-source specific state (*, G).
Extensive	Upstream state—Information about upstream interface.

Output Fields

Group—Group address.

Source—Multicast source.

RP—Rendezvous point for the PIM group.

Flags—PIM flags. It can be one or more of the following:

- dense—Dense mode entry.
- rptree—Entry is on the Rendezvous Point tree.
- sparse—Sparse mode entry.
- spt—Entry is on the shortest-path tree for the source.
- spt-pending—The software has initiated the switch to the shortest-path tree for this source.
- wildcard—Indicates that this entry is on the shared tree.

Upstream interface—RPF Interface toward the source address for the source specific state (S, G) or toward the RP address for the non-source specific state (*, G).

Upstream state—(Extensive output only) Information about the upstream interface. It can be zero or more of the following:

- Join to RP—Sending a join to the Rendezvous Point.
- Join to Source—Sending a join to the Source.
- Local RP—This router is the RP and hence is sending neither joins nor prunes toward the RP.
- Local Source—Source is locally attached to this router and hence is sending neither joins nor prunes toward the source.
- Prune to RP—Sending a prune to the Rendezvous Point.
- Prune to Source—Sending a prune to the source.

Downstream interfaces—(Detail and Extensive output only) Detailed information about downstream interfaces, including interface name, state, flags, and timeout.

Downstream neighbors—(Extensive output only) Detailed information about the downstream interfaces:

- Interface address—Address of downstream neighbor.
- State—Information about the downstream neighbor. Can be join or prune.
- Flags—PIM join flags. It can be zero or more of the following: S (Sparse), R (RPtree), and W (Wildcard).

Timeout—(Extensive output only) Time remaining until the neighbor is declared to be inactive, in seconds. If there is a directly connected host, the timeout is Infinity.

Sample Output: show pim join brief

```
user@host> show pim join brief
show pim join brief
Instance: PIM.master

Group           Source          RP              Flags
239.1.1.1       0.0.0.0         10.255.71.47    sparse,rptree,wildcard
    Upstream interface: so-0/0/2.0
239.1.1.1       10.255.71.46                    sparse
    Upstream interface: local
```

Sample Output: show pim join (instance) brief

```
user@host> show pim join instance VPN-A brief
show pim join brief instance VPN-A
Instance: PIM.VPN-A

Group           Source          RP              Flags
225.1.1.1       0.0.0.0         10.255.245.91   sparse,rptree,wildcard
    Upstream interface: mt-1/1/0.32769
```

Sample Output: show pim join detail

```
user@host> show pim join detail
show pim join detail
Instance: PIM.master

Group           Source          RP              Flags
239.1.1.1       0.0.0.0         10.255.71.47    sparse,rptree,wildcard
    Upstream interface: so-0/0/2.0
    Downstream interfaces:
    mt-1/1/0.32769
239.1.1.1       10.255.71.46                    sparse
    Upstream interface: local
    Downstream interfaces:
    so-0/0/2.0
```

Sample Output: show pim join extensive

```
user@host> show pim join extensive
Instance: PIM.master

Group           Source          RP              Flags
239.1.1.1       0.0.0.0         10.255.71.47    sparse,rptree,wildcard
    Upstream interface: so-0/0/2.0
    Upstream State: Join to RP
    Downstream Neighbors:
    Interface: mt-1/1/0.32769
        10.255.71.46     State: Join   Flags: SRW  Timeout: Infinity

239.1.1.1       10.255.71.46                    sparse
    Upstream interface: local
    Upstream State: Local Source, Prune to RP
    Downstream Neighbors:
    Interface: so-0/0/2.0
        192.168.196.70   State: Join   Flags: S    Timeout: 166
```

Sample Output: show pim join (instance) extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A

Group           Source          RP              Flags
225.1.1.1       0.0.0.0         10.255.245.91   sparse,rptree,wildcard
    Upstream interface: mt-1/1/0.32769
    Upstream State: Join to RP
    Downstream Neighbors:
    Interface: t1-1/0/0:0.0
        192.168.196.106  State: Join   Flags: SRW  Timeout: 207
```

show pim neighbors

Syntax show pim neighbors <brief | detail> < *instance* >

Description Display information about PIM neighbors.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

instance—(Optional) Display all interfaces under the named routing instance.

Default: brief

Required Privilege Level view

Output Fields Instance—Name of the routing instance.

Interface—Interface through which the neighbor is reachable.

DR priority—For PIM Version 2 only, the interface's priority to become the designated router.

Neighbor addr—Address of the neighboring PIM router.

V—PIM version running on the neighbor. It can be either 1 or 2.

Mode—Mode. It can be Sparse, Dense, SparseDense, or Unknown. When the neighbor is running PIM Version 2, this mode is always Unknown.

Holdtime—Lifetime of the neighbor, in seconds.

Timeout—How long until the PIM neighbor is declared to be inactive, in seconds.

Sample Output: show PIM neighbor for master (default) instance

```
user@host> show pim neighbors
Instance: PIM.master

Interface        DR priority Neighbor addr   V Mode        Holdtime Timeout
so-0/0/2.0                1 192.168.196.70  2 Unknown          105      93
```

Sample Output: show PIM neighbor for specific PIM instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A

Interface        DR priority Neighbor addr   V Mode        Holdtime Timeout
t1-1/0/0:0.0              1 192.168.196.106 2 Unknown          105      79
mt-1/1/0.32769            1 10.255.71.51    2 Unknown          105      88
```

show pim rps

Syntax show pim rps <*group-address*> <brief | detail | extensive><*instance*>

Description Display information about PIM rendezvous points (RPs).

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information, including RP calculations.

extensive—(Optional) Display very detailed information.

group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists which router should be the RP for that group.

instance—(Optional) Display all interfaces under the named routing instance.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show pim rps on page 684
Sample Output: show pim rps detail on page 684

Options at a Glance Table 86 summarizes which information is included in each show pim rps command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 86: Show PIM RPS Output Fields Summary

Options	Field Description
All	Active groups—Number of groups currently using this RP.
Brief	Group prefixes—Addresses of groups that this RP spans.
Detail	Group Ranges—Addresses of groups that this RP spans.
All	Holdtime—How long to keep the RP active, in seconds.
All	Instance—Name of the routing instance.
Detail	Learned—Address and method by which the RP was learned.
All	RP address—Address of the RP.
Detail	Time Active—How long the RP has been active.
Brief	Timeout—How long until the local router declares the RP to be unreachable, in seconds.
Detail	Total—Total number of active groups for this RP.
Brief	Type—Type of RP.

Output Fields

Instance—Name of the routing instance.

RP address—Address of the RP.

Learned—(Detail output only) Address and method by which the RP was learned.

Time active—(Detail output only) How long the RP has been active.

Type—(Brief output only) Type of RP. It can be auto-rp, bootstrap, or static.

Holdtime—How long to keep the RP active, in seconds.

Group ranges—(Detail output only) Addresses of groups that this RP spans.

Timeout—(Brief output only) How long until the local router declares the RP to be unreachable, in seconds.

Active groups—Number of groups currently using this RP.

Group prefixes—(Brief output only) Addresses of groups that this RP can span.

Total—(Detail output only) Total number of active groups for this RP.

Sample Output: show pim rps

```
user@host> show pim rps
show pim rps
Instance: PIM.master

RP address      Type      Holdtime Timeout Active groups Group prefixes
10.255.71.47    static           0    None             1 224.0.0.0/4
```

Samplle Output: show pim rps (instance)

```
user@host> show pim rps instance VPN-A
show pim rps instance VPN-A
Instance: PIM.VPN-A

RP address      Type      Holdtime Timeout Active groups Group prefixes
10.255.245.91   static           0    None             1 224.0.0.0/4
```

Sample Output: show pim rps detail

```
user@host> show pim rps detail
show pim rps detail
Instance: PIM.master

RP: 10.255.71.47
Learned via: static configuration
Time Active: 07:22:55
Holdtime: 0
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    239.1.1.1

    total 1 groups active
```

Sample Output: show pim rps detail (instance)

```
user@host> show pim rps detail instance VPN-A
show pim rps detail instance VPN-A
Instance: PIM.VPN-A

RP: 10.255.245.91
Learned via: static configuration
Time Active: 07:22:59
Holdtime: 0
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    225.1.1.1

    total 1 groups active
```

Sample Output: show pim rps extensive (instance)

```
user@host> show pim rps extensive instance VPN-A
show pim rps extensive instance VPN-A
Instance: PIM.VPN-A

RP: 10.255.245.91
Learned via: static configuration
Time Active: 07:22:16
Holdtime: 0
Device Index: 33
Subunit: 32770
Interface: pe-1/1/0.32770
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    225.1.1.1

    total 1 groups active

Register State for RP:

Group           Source          FirstHop        RP Address      State    Timeout
```

show pim source

Syntax show pim source <*source-prefix*> <brief | detail> < *instance* >

Description Display information about the PIM source RPF state.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information, including RP calculations.

source-prefix—(Optional) Display the state for source RPFs in the given range.

instance—(Optional) Display all interfaces under the named routing instance.

Default: brief

Required Privilege Level view

Output Fields Instance—Name of the routing instance.

RPF Address—Address of source or reverse-path.

Prefix/length—Prefix and prefix length for the route used to reach the Reverse-path-forwarding address.

Upstream interface—RPF Interface toward the source address.

Neighbor address—Address of RPF neighbor used to reach the source address.

Sample Output: show pim source (standard, instance)

```
user@host> show pim source VPN-A
show pim source VPN-A
Instance:PIM.VPN-A
RPF Address     Prefix/length      Upstream interface     Neighbor address
10.10.10.101    10.10.10.101   /32 local                  Local
10.10.10.103    10.10.10.103   /32 so-0/1/0.0             10.10.20.10
```

Sample Output: show pim source detail

```
user@host> show pim source detail
show pim source detail
Instance: PIM.master

RPF Address     Prefix/length      Upstream interface     Neighbor address
10.255.71.46    10.255.71.46/32    local                  Local
(10.255.71.46, 239.1.1.1)
10.255.71.47    10.255.71.47/32    so-0/0/2.0             192.168.196.70
(*, 239.1.1.1)
```

Sample Output: show pim source detail (instance)

```
user@host> show pim source detail instance VPN-A
show pim source detail instance VPN-A
Instance: PIM.VPN-A

RPF Address     Prefix/length      Upstream interface     Neighbor address
10.255.245.91   10.255.245.91/32   mt-1/1/0.32769         10.255.71.51
(*, 225.1.1.1)
```

show pim statistics

Syntax show pim statistics <interface *interface-name*> < *instance* >

Description Display PIM statistics.

Options interface *interface-name*—(Optional) Display statistics about the specified interface. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

instance—(Optional) Display all interfaces under the named routing instance.

Required Privilege Level view

See Also clear pim statistics on page 646

Output Fields Instance—Name of the routing instance.

PIM statistics on all interfaces or PIM statistics for interfaces *interface-name*—Name of the interface for which the statistics are being reported.

Received—Number of received statistics.

Sent—Number of messages sent of a certain type.

Rx errors—Number of received packets that contained errors.

PIM statistics summary for all interfaces—Summary of PIM statistics for all interfaces.

Sample Output

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                   2102       4203          0
V1 Register                   0          0          0
V1 Register Stop              0          0          0
V1 Join Prune             14153      13074          0
V1 RP Reachability            0          0          0
V1 Assert                     0          0          0
V1 Graft                      0          0          0
V1 Graft Ack                  0          0          0

PIM statistics summary for all interfaces:
Unknown type                             0
V1 Unknown type                          0
Unknown Version                          0
Neighbor unknown                         0
Bad Length                               0
Bad Checksum                             0
Bad Receive If                           0
Rx Intf disabled                      1998
Rx V1 Require V2                         0
```

```
Rx Register not RP                      0
RP Filtered Source                      0
Unknown Reg Stop                        0
Rx Join/Prune no state               1034
Rx Graft/Graft Ack no state             0
Rx Graft on upstream if                 0
Rx CRP not BSR                          0
Rx BSR when BSR                         0
Rx BSR not RPF if                       0
Rx unknown hello opt                    0
Rx data no state                        0
Rx RP no state                          0
Rx aggregate                            0
Rx malformed packet                     0
No RP                                   0
No route upstream                       0
RP mismatch                             0
RPF neighbor unknown                    0
```

show sap listen

Syntax show sap listen <brief | detail>

Description Display the addresses that the router is listening to so that it receives multicast session announcements.

Options brief—(Optional) Display brief information.

detail—(Optional) Display detailed information.

Default: brief

Required Privilege Level view

Output Fields Group address—Address of the group that the local router is listening to for SAP messages.

Port—UDP Port number used for sap.

Sample Output

```
user@host> show sap listen
Group address   Port
224.2.127.254   9875
239.255.255.255 9875
```

test msdp

Syntax test msdp <dependent-peers *prefix* | rpf-peer *originator*>

Description Find MSDP peers.

Options dependent-peers *prefix*—(Optional) Display downstream dependent MSDP peers.

rpf-peer *originator*—(Optional) Display the MSDP reverse-path-forwarding peer for the originator.

Required Privilege Level view

Part 7
MPLS Applications

- MPLS Monitoring and Troubleshooting on page 691
- RSVP Monitoring and Troubleshooting on page 717
- LDP Monitoring and Troubleshooting on page 737
- Layer 2 VPN Monitoring and Troubleshooting on page 751
- Layer 3 VPN Monitoring and Troubleshooting on page 779

Chapter 42
MPLS Monitoring and Troubleshooting

Table 87 summarizes the command-line interface (CLI) commands you can use to monitor MPLS. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 87: Commands for Monitoring MPLS

Task Category	Task or Information to Monitor	Command
MPLS interfaces	Status of interfaces on which MPLS is running	show mpls interface on page 698
Named paths	Configured named paths that are used in dynamic MPLS	show mpls path on page 709
GMPLS	Link-management peers and resources.	show link-management on page 695
LSPs	Configured LSPs on this router, as well as all ingress, transit, and egress LSPs	show mpls lsp on page 699
	Disconnect and restart dynamic LSPs that originate from this router	clear mpls lsp on page 692
MPLS administrative groups	MPLS administrative groups	show mpls admin-groups on page 696
CSPF statistics	CSPF statistics	show mpls cspf on page 697
Traffic engineering database	Entries in the TED database	show ted database on page 710
	Protocols contributing to the TED	show ted protocol on page 716
	Current TED links	show ted link on page 714
Circuit cross-connects	Information about configured cross-connects	show connections on page 693

clear mpls lsp

Syntax clear mpls lsp <autobandwidth> <name *name*> <optimize | optimize-aggressive> < path > <statistics>

Description Disconnect and restart LSPs. This operation disconnects existing RSVP sessions on the ingress router, releases the routes and states associated with the LSPs, and then starts new LSPs. Issuing this command might impact traffic travelling along the LSPs, because there might be a time lag between the old path being torn down and the new path being set up.

Options none—Reset and restart *all* LSPs that originated from this router; that is, all LSPs for which this router is the ingress router. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in "Wildcard Characters in Interface Names" on page 187.

optimize—(Optional) Run nonpreemptive optimization computation now.

optimize-aggressive—(Optional) Run aggressive optimization computation now.

path—(Optional) Clear specific LSP path matching pattern.

statistics—(Optional) Clear LSP statistics.

Required Privilege Level clear

See Also show rsvp session on page 727

Sample Output

```
user@host> clear mpls lsp
user@host>
```

show connections

Syntax show connections <all | interface-switch | lsp-switch | remote-interface-switch>
<up | down | up-down> <brief | extensive | history | labels | status> <*name*>

Description Display information about the configured CCC connections.

Options all—(Optional) Display all connections.

brief—(Optional) Display brief connection status information.

down—(Optional) Display connections that are not operational.

extensive—(Optional) Display very detailed connection, labels, and history information.

history—(Optional) Display connection history.

interface-switch—(Optional) Display interface switch connections.

labels—(Optional) Display labels used for transmit and receive LSPs.

lsp-switch—(Optional) Display LSP switch connections.

name—(Optional) Display information about the named connection.

remote-interface-switch—(Optional) Display remote interface switch connections.

status—(Optional) Display connection and circuit status.

up—(Optional) Display connections that are operational.

up-down—(Optional) Display both operational and nonoperational connections.

Default: brief, status

Required Privilege Level view

Output Fields Connection/Circuit—Name of the configured CCC connection.

Type—Type of connection or circuit:

- intf—Interface circuit.
- if-sw—Layer 2 switching cross-connect.
- lsp-sw—LSP stitching cross-connect.
- rlsp—Receive LSP circuit.
- rmt-if—Remote interface cross-connect.
- tlsp—Transmit LSP circuit.

Stat—Connection or circuit status. Refer to the output's legend for an explanation of the status field values.

Time last up—Time that the connection or circuit last transitioned to the Up (operational) state.

Up trans—Number of times that the connection or circuit has transitioned to the Up (operational) state.

Sample Output

```
user@host> show connections all
CCC connections
  Legend for status (St)              Legend for connection types
  UN -- uninitialized                 if-sw:  interface switching
  NP -- not present                   rmt-if: remote interface switching
  WE -- wrong encapsulation           lsp-sw: LSP switching
  DS -- disabled
  Dn -- down                          Legend for circuit types
  -> -- only outbound conn is up      intf -- interface
  <- -- only inbound  conn is up      tlsp -- transmit LSP
  Up -- operational                   rlsp -- receive LSP

Connection/Circuit                  Type    St  Time last up     # Up trans
intf-sw                             if-sw   Dn
  at-1/2/3.4                          intf  NP
  at-3/2/1.0                          intf  NP
pr-pr                               lsp-sw  Dn
  from-tuscon                         rlsp  NP
  to-tuscin                           tlsp  Dn
rmt-intf-sw                         rmt-if  Dn
  so-5/3/1.0                          intf  NP
  blue23                              tlsp  Dn
  green12                             rlsp  NP
```

show link-management

Syntax show link-management <name *name*> <peer | te-link> <routing>

Description Displays link-management peers and resources.

Required Privilege Level view

Options name—Name of the peer or TE link to be displayed.

peer—Name of the neighbor that shares a forwarding adjacency.

routing—Displays link-management peers and resources from the routing daemon rather than the link-management daemon.

te-link—Logical collection of resources used to set up traffic-engineered forwarding paths.

Output Fields Peer Name—Name of the peer.

System identifier—Internal identifier for the peer. The range of values can be 0-64,000.

State—State of the peer or TE link. It can be Up or Down.

Control address—Address to which a control channel is established.

Control interface—Interface over which control packets are sent.

TE links—Traffic-engineered links that are managed by the peer they are associated with.

Local identifier— Identifier of the local side of the link.

Remote identifier— Identifier of the remote side of the link.

Local address—Address of the local side of the link.

Remote address—Address of the remote side of the link.

Encoding—Describes the physical layer media type determined by the interfaces contained in the TE link. Typical values include SDH/SONET, Ethernet, and PDH.

Minimum bandwidth—Smallest single allocation of bandwidth possible on the TE link. This number is equal to the smallest bandwidth interface that is a member of the TE link (in bps).

Maximum bandwidth—Largest single allocation of bandwidth possible on the TE link. This number is equal to the largest bandwidth interface that is a member of the TE link (in bps).

Total bandwidth—Sum of the bandwidths of all interfaces that are members of the TE link (in bps).

Available Bandwidth—Sum of the bandwidths of all interfaces that are members of the TE link and that are not yet allocated (in bps).

Name—Name of the interface.

Local ID—Identifier of the local side of the interface.

Remote ID—Identifier of the remote side of the interface.

Bandwidth—Bandwidth of the member interface (in bps).

In use—Whether the resource is allocated to an LSP. It can be Yes or No.

Sample Output

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11968
 State: Up, Control address: 10.255.245.194, Control interface: fe-0/0/0.0
  TE links:
   TElnk2
   TElnk1

TE link name: TElnkDC1, State: Up
  Local identifier: 4638, Remote identifier: 3837, Local address: 1.1.1.197,
  Remote address: 1.1.1.196, Encoding: Ethernet, Minimum bandwidth: 100000000,
  Maximum bandwidth: 1000000000, Total bandwidth: 1100000000,
  Available bandwidth: 1000000000
  Name          Local ID  Remote ID   Bandwidth In use
    ge-1/0/0         21269      21269  1000000000 No
    fe-2/0/0         46371      46371   100000000 Yes
```

show mpls admin-groups

Syntax show mpls admin-groups

Description Display information about the configured MPLS administrative groups.

Required Privilege Level view

Output Fields Group—Name of the administrative group.

Bit index—Value assigned to the administrative group.

Sample Output

```
[edit]
user@host# show
mpls {
    admin-groups{
        green 0;
        gold 1;
        blue 2;
        black 3;
    }
...
[edit]
user@host# exit
user@host> show mpls admin-groups
Group        Bit index
black         3
blue          2
gold          1
green         0
```

show mpls cspf

Syntax show mpls cspf

Description Display CSPF statistics.

Required Privilege Level view

Output Fields Queue length—Number of LSPs queue for automatic path computation.

Current—Current queue length.

Maximum—Maximum queue length (high-water mark).

Dequeued—Number of aborted computation attempts.

Paths—Counters for label-switched path computations.

Total—Sum of the next four fields.

Successful—Number of path computations that completed successfully.

No Route—Number of path computations that failed because the destination is unreachable.

Sys Error—Number of path computations that failed because of lack of memory.

CSPFs—Total number of CSPF computations. A single path might require multiple CSPF computations.

Time—Time required to perform the label-switched path computation, in sections.

Total—Total amount of time consumed by the CSPF path computation algorithm.

CSPFs—Total number of CSPF computations.

Avg per CSPF—Average amount of time required for each CSPF computation.

% of rpd—Percentage of routing process CPU used in the CSPF computation.

Sample Output

```
user@host> show mpls cspf
CSPF statistics
Queue length   current       maximum      dequeued
                     0             0             0

Paths            total    successful      no route     sys error         CSPFs
                     0             0             0             0             0

Time (secs)      total         CSPFs  avg per CSPF      % of rpd
              0.000000      0.000000      0.000000        0.0000
```

show mpls interface

Syntax show mpls interface

Description Display information about interfaces on which MPLS is enabled. MPLS is enabled on an interface when the interface is configured in both the mpls and interface sections of the configuration hierarchy (with the set protocol mpls interface *interface-name* and set interface *interface-name* unit 0 family mpls statements, respectively).

Required Privilege Level view

Output Fields Interface—Name of the interface.

State—State of the interface. It can be Up or Down.

Administrative groups—The administratively assigned colors of the link.

Sample Output

```
user@host> show mpls interface
Interface   State        Administrative groups
so-1/0/0.0  Up           Blue Yellow Red
```

show mpls lsp

Syntax show mpls lsp <name *name*> <*lsp-type*> <statistics> <brief | terse | detail | extensive> <up | down>

Description Display information about configured and active dynamic LSPs in which this router participates.

Options brief—(Optional) Display brief LSP information.

detail—(Optional) Display detailed LSP information, including all current state information and all parameters and counters that are part of the LSP.

down—(Optional) Display only LSPs that are inactive.

extensive—(Optional) Display all information about LSPs, including all past state history and the reasons why an LSP failed. The show mpls lsp extensive command displays the same information as the show mpls lsp detail command but covers the most recent 50 events.

> **Note** CSPF always results in a strictly routed LSP, even if the configuration specifies a loose route, or no route at all.

lsp-type—(Optional) Display information about a particular LSP type:

- egress—Sessions that terminate on this router.
- ingress—Sessions that originate from this router.
- transit—Sessions that transit through this router.

The output is displayed in the brief format by default.

The output of the show mpls lsp egress and show mpls lsp transit commands are identical to that of the show rsvp session egress lsp and clear rsvp session transit lsp commands, respectively, because RSVP is the signaling protocol that is used to set up LSPs.

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the LSP name, as described in “Wildcard Characters in Interface Names” on page 187.

statistics—(Optional) Display accounting information about LSPs. This option is valid only on the ingress and on transit routers, because only these routers can maintain statistics. No statistics are available for LSPs on the egress router, because the penultimate router in the LSP sets the label to 0. Also, as the packet arrives at the egress router, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

terse—(Optional) Display only a count of the number of sessions and the number of up and down LSPs. Specifying this option is useful when there are a large number of sessions.

up—(Optional) Display only LSPs that are active.

Default: brief; display active and inactive LSPs

Required Privilege Level view

Sample Output

Sample Output: show mpls lsp terse on page 706
Sample Output: show mpls lsp (standard, brief) on page 706
Sample Output: show mpls lsp detail on page 706
Sample Output: show mpls lsp extensive on page 708

Options at a Glance Table 88 summarizes which information is included in each show mpls lsp command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 88: Show MPLS LSP Output Field Summary

Options	Field Description
Detail Extensive	Active path indicator—(Ingress LSP only) A value of * indicates that the path is active. The absence of * indicates that the path is not active.
Detail Extensive	Active Route—Number of active routes (prefixes) that have been installed in the forwarding table.
Detail Extensive	address—(First line of each section) Destination (egress router) of the LSP.
Detail Extensive	AdjustTimer—(Ingress LSP only) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.
Detail Extensive	Autobandwidth—(Ingress LSP only) The LSP is performing autobandwidth allocation.
Detail Extensive	Bandwidth—(Ingress LSP only) Active bandwidth for the LSP path, in bps.
Detail Extensive	Bandwidth Adjustment in xxx second(s)—(Ingress LSP only) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.
Brief Standard	Bidir—The LSP allows data to travel in both directions between GMPLS devices.
Detail Extensive	Bidirectional—The LSP allows data to travel both ways between GMPLS devices.
Detail Extensive	Computed ERO (S [L] denotes strict [loose] hops)—(Ingress LSP only) Computed Explicit Route is a series hops. Each hop has an address followed by a hop indicator. The value of the hop indicator can be strict S or loose L.
Detail Extensive	COS—(Ingress LSP only) Class-of-Service value.
Extensive	Created—(Ingress LSP only) Date and time the LSP was created, displayed as day of the week *abc*, month *abc*, day *xx*, time in *hours:minutes:seconds*, and year *xxxx*.
Detail Extensive	CSPF metric—(Ingress LSP only) Constrained Shortest Path First metric for this path.
Extensive	Date—(Ingress LSP only) Date of the LSP event, displayed as month *abc* and day *xx*.
Brief Detail Extensive Terse	Egress RSVP—Information about the LSPs on the egress router. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.
Detail Extensive	Encoding type—Describes the nature of the LSP. The encoding type for the LSP can be Ethernet, PDH, or SDH/SONET.
Extensive	Event—(Ingress LSP only) Description of the LSP event.
Brief Detail Extensive	From—Source (ingress router) of the session.
Detail Extensive	GPID—Generalized Payload Identifier (identifier of the payload carried by an LSP). It can be HDLC, Ethernet, IPv4, PPP, or Unknown.

Options	Field Description
Extensive	Index number—(Ingress LSP only) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.
All	Ingress LSP—Information about LSPs on the ingress router. Each session has one line of output.
Brief Detail Extensive	Labelin—Incoming label for this LSP.
Brief Detail Extensive	Labelout—Outgoing label for this LSP.
Detail Extensive	LoadBalance—(Ingress LSP only) Indicates which CSPF load-balancing rule was configured to select the LSP's path among equal-cost paths. It can be Most-fill, Least-fill, or Random.
Brief Detail Extensive	LSP name—Name of the LSP.
Brief Detail Extensive	LSP state—State of the LSP that is being handled by this RSVP session. It can be either Up or Dn (down).
Detail Extensive	MaxAvgBW util—(Ingress LSP only) Current value of the actual maximum average bandwidth utilization, in bps.
Detail Extensive	MaxBW—(Ingress LSP only) Configured maximum value of the LSP, in bps.
Detail Extensive	MinBW—(Detail and extensive output only, Ingress LSP only) Configured minimum value of the LSP, in bps.
Detail Extensive	OptimizeTimer—(Ingress LSP only) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.
Detail Extensive	PATH rcvfrom—Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor.
Detail Extensive	PATH sentto—Address of the next-hop (downstream) router or client, interface used to reach this neighbor, and how many packets were sent to the downstream router.
Detail Extensive	Port number—Protocol ID and sender or receiver port used in this RSVP session.
Detail Extensive	Primary—(Ingress LSP only) Name of the primary path.
Detail Extensive	Protection—Describes the configured protection capability desired for the LSP. It can be Extra, Enhanced, none, One plus one, One to one, or Shared.
Detail Extensive	Received RRO—(Ingress LSP only) Received Record Route is a series of hops. Each hop has an address followed by a flag.
Detail Extensive	Record route—Recorded route for the session as taken from the Record Route Object.
Detail Extensive	Reoptimization in xxx seconds—(Ingress LSP only) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.
Detail Extensive	RESV rcvfrom—Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.
Brief Detail Extensive	Resv style—RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).
Brief	Rt—Number of active routes (prefixes) that have been installed in the routing table.
Detail Extensive	Signal type—Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP, and can have values of DS0, DS3, STS-1, STM-1, or STM-4.
Detail Extensive	Since—Date and time the RSVP session was initiated.
Detail Extensive	Standby—(Ingress LSP only) Name of the path in standby mode.
All	State—State of the path. It can be Up or Dn (down).
Detail Extensive	Suggested label receive—Label the upstream node suggests to use in the Resv message that is sent. Displayed only when the LSP is bidirectional.
Detail Extensive	Suggested label sent—Label the downstream node suggests to use in its Resv message that is returned. Displayed only when the LSP is bidirectional.

Options	Field Description
Detail Extensive	Switching type—Describes the type of switching on the links needed for the LSP. It can be Fiber, Lamda, Packet, TDM, or PSC-1.
Extensive	Time—(Ingress LSP only) Time of the LSP event in displayed as *hours:minutes:seconds*.
Detail Extensive	Time left—Number of seconds remaining in the lifetime of the reservation.
Brief	To—Destination (egress router) of the session.
Brief Detail Extensive Terse	Transit RSVP—Number of LSPs on the transit routers and the state of these paths.
Detail Extensive	Tspec—Sender's traffic specification, which describes the sender's traffic parameters.
Detail Extensive	Upstream label in—Incoming label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.
Detail Extensive	Upstream label out—Outgoing label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Output Fields

Ingress LSP—Information about LSPs on the ingress router. Each session has one line of output.

Egress RSVP—Information about the LSPs on the egress router. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.

Transit RSVP—Number of LSPs on the transit routers and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.

address—(Detail output only; first line of each section) Destination (egress router) of the LSP.

To—(Brief output only) Destination (egress router) of the session.

From—(Brief and detail output only) Source (ingress router) of the session.

LSP state—(Brief and detail output only) State of the LSP handled by this RSVP session. It can be either Up or Dn (down).

Active Route—(Detail and extensive output only) Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).

LSPname—(Brief and detail output only) Name of the LSP.

Bidir—The LSP allows data to travel in both directions between GMPLS devices.

Bidirectional—The LSP allows data to travel both ways between GMPLS devices.

Rt—(Brief output only) Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).

ActivePath—(Detail and extensive output only, ingress LSP only) Name of the active path. It can be Primary or Secondary.

LoadBalance—(Detail and extensive output only, ingress LSP only) Indicates which CSPF load-balancing rule was configured to select the LSP's path among equal-cost paths. It can be Most-fill, Least-fill, or Random.

Signal type—Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP, and can have values of DS0, DS3, STS-1, STM-1, or STM-4. The default is none.

Encoding type—Describes the nature of the LSP. The encoding type for the LSP can be Ethernet, PDH, or SDH/SONET.

Switching type—Describes the type of switching on the links needed for the LSP. It can be Fiber, Lamda, Packet, TDM, or PSC-1.

GPID—Generalized Payload Identifier (identifier of the payload carried by an LSP). It can be HDLC, Ethernet, IPv4, PPP, or Unknown.

Protection—Describes the configured protection capability desired for the LSP. It can be Extra, Enhanced, none, One plus one, One to one, or Shared.

Upstream label in—Incoming label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Upstream label out—Outgoing label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Suggested label received—Label the upstream node suggests to use in the Resv message that is sent. Displayed only when the LSP is bidirectional.

Suggested label sent—Label the downstream node suggests to use in the Resv message that is returned. Displayed only when the LSP is bidirectional.

Autobandwidth—(Detail and extensive output only, ingress LSP only) The LSP is performing autobandwidth allocation.

MinBW—(Detail and extensive output only, ingress LSP only) Configured minimum value of the LSP, in bps.

MaxBW—(Detail and Extensive output only, ingress LSP only) Configured maximum value of the LSP, in bps.

AdjustTimer—(Detail and extensive output only, ingress LSP only) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.

MaxAvgBW util—(Detail and extensive output only, ingress LSP only) Current value of the actual maximum average bandwidth utilization, in bps.

Bandwidth Adjustment in xxx second(s)—(Detail and extensive output only, ingress LSP only) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.

Active path indicator—(Detail and extensive output only, ingress LSP only) A value of * indicates that the path is active. The absence of * indicates that the path is not active.

In the following example, “long” is the active path.

```
*Primary   long
 Standby   short
```

Primary—(Detail and extensive output only, ingress LSP only) Name of the primary path.

Secondary—(Detail and extensive output only, ingress LSP only) Name of the secondary path.

Standby—(Detail and extensive output only, ingress LSP only) Name of the path in standby mode.

State—(Detail and extensive output only, ingress LSP only) State of the path can be Up or Dn (down).

COS—(Detail and extensive output only, ingress LSP only) Class-of-Service value.

Bandwidth—(Detail and extensive output only, ingress LSP only) Active bandwidth for the LSP path, in bps.

OptimizeTimer—(Detail and extensive output only, ingress LSP only) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.

Reoptimization in xxx seconds—(Detail and extensive output only, ingress LSP only) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.

Computed ERO (S [L] denotes strict [loose] hops)—(Detail and extensive output only, ingress LSP only) Computed Explicit Route is a series hops. Each hop has an address followed by a hop indicator. The value of the hop indicator can be strict S or loose L.

CSPF metric—(Detail and extensive output only, ingress LSP only) Constrained Shortest Path First metric for this path.

Received RRO—(Detail and extensive output only, ingress LSP only) Received Record Route is a series of hops. Each hop has an address followed by a flag. (In most cases, the Received Record Route is the same as the Computed Explicit Route. If the Received RRO is different from the Computed ERO, there is a toplogy change in the network, and the route is taking a detour.) Some possible flag values are:

- 0x01—Local protection available. The link downstream of this node is protected by a local repair mechanism. This flag can only be set if the Local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding Path message.
- 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).
- 0x03—Combination of 0x01 and 0x02.
- 0x04— Bandwidth protection. The downstream router has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.

- 0x08—Node protection. The downstream router has a backup path providing protection against link and node failure on the corresponding path section. If the downstream router could only set up a link-protection backup path, the "Local protection available" bit will be set but the "Node protection" bit is cleared.
- 0x09—Detour is established. Combination of 0x01 and 0x08.
- 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08.

Index number—(Extensive output only, ingress LSP only) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.

Date—(Extensive output only, ingress LSP only) Date of the LSP event, displayed as month *abc* and day *xx*.

Time—(Extensive output only, ingress LSP only) Time of the LSP event in displayed as *hours:minutes:seconds*.

Event—(Extensive output only, ingress LSP only) Description of the LSP event.

Created—(Extensive output only, ingress LSP only) Date and time the LSP was created, displayed as day of the week *abc*, month *abc*, day *xx*, time in *hours:minutes:seconds*, and year *xxxx*.

Resv style—(Brief and detail output only) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).

Labelin—(Brief and detail output only) Incoming label for this LSP.

Labelout—(Brief and detail output only) Outgoing label for this LSP.

LSPname—(Brief and detail output only) Name of the LSP.

Time left—(Detail output only) Number of seconds remaining in the lifetime of the reservation.

Since—(Detail output only) Date and time when the RSVP session was initiated.

Tspec—(Detail output only) Sender's traffic specification, which describes the sender's traffic parameters.

Port number—(Detail output only) Protocol ID and sender or receiver port used in this RSVP session.

PATH rcvfrom—(Detail output only) Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor.

PATH sentto—(Detail output only) Address of the next-hop (downstream) router or client, interface used to reach this neighbor, and how many packets were sent to the downstream router.

RESV rcvfrom—(Detail output only) Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.

Record route—(Detail output only) Recorded route for the session as taken from the Record Route Object.

Sample Output: show mpls lsp terse

```
user@host> show mpls lsp terse

Ingress LSP: 2 label-switched paths
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
Total 1 displayed, Up 1, Down 0
```

Sample Output: show mpls lsp (standard, brief)

```
user@host> show mpls lsp
Ingress LSP: 2 sessions
To              From            State Rt ActivePath       P     LSPname
10.255.245.196  10.255.245.194  Up     1                  *     pro3-ac
10.255.245.196  10.255.245.194  Up     0 prim1            *     Glsp1 Bidir
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

Sample Output: show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 2 sessions

10.255.245.196
  From: 10.255.245.194, State: Up, ActiveRoute: 1, LSPname: pro3-ac
  ActivePath:  (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Protection: None
 *Primary                    State: Up
    Bandwidth: 20Mbps
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 40)
          192.168.207.30 S 192.168.207.38 S
    Received RRO:
          192.168.207.30 192.168.207.38
```

```
10.255.245.196
  From: 10.255.245.194, State: Up, ActiveRoute: 0, LSPname: Glsp1
  Bidirectional
  ActivePath: prim1 (primary)
  LoadBalance: Random
  Signal type: STM-1
  Encoding type: Fiber, Switching type: Fiber, GPID: PPP
  Protection: None
 *Primary   prim1            State: Up
    Received RRO:
          2.2.2.2
Total 2 displayed, Up 2, Down 0

Egress LSP: 2 sessions

10.255.245.194
  From: 10.255.245.195, LSPstate: Up, ActiveRoute: 0
  LSPname: Gpro3-ba
  Bidirectional: Upstream label in: -, Upstream label out: 39811
  Suggested label received: 39811, Suggested label sent: -
  Resv style: 1 FF, Label in: 39811, Label out: -
  Time left:  123,  Since: Mon Apr  8 12:11:38 2002
  Tspec: rate 0bps size 0bps peak 622.08Mbps m 20 M 1500
  Port number: sender 1 receiver 17 protocol 0
  PATH rcvfrom: 192.168.207.132  (fe-0/0/0.0) 9 pkts
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 4.4.4.2  <self>

10.255.245.194
  From: 10.255.245.195, LSPstate: Up, ActiveRoute: 0
  LSPname: pro3-ba
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  118,  Since: Mon Apr  8 12:12:43 2002
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 19 protocol 0
  PATH rcvfrom: 192.168.207.26  (so-1/0/0.0) 9 pkts
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 192.168.207.26  <self>
Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions

10.255.245.198
  From: 10.255.245.197, LSPstate: Up, ActiveRoute: 0
  LSPname: pro3-de
  Resv style: 1 SE, Label in: 100000, Label out: 3
  Time left:  131,  Since: Mon Apr  8 12:12:56 2002
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 16 protocol 0
  FastReroute desired
  PATH rcvfrom: 192.168.207.130  (fe-0/0/0.0) 8 pkts
  PATH sentto: 192.168.207.30 (so-0/1/1.0) 9 pkts
  RESV rcvfrom: 192.168.207.30  (so-0/1/1.0) 9 pkts
  Explct route: 192.168.207.30
  Record route: 192.168.207.130  <self>  192.168.207.30
    Detour is Up
    Detour PATH sentto: 192.168.207.34 (so-0/1/0.0) 8 pkts
    Detour RESV rcvfrom: 192.168.207.34  (so-0/1/0.0) 8 pkts
    Detour Explct route: 192.168.207.34
    Detour Record route: 192.168.207.130  <self>  192.168.207.34
Total 1 displayed, Up 1, Down 0
```

Sample Output: show mpls lsp extensive

```
user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

10.255.245.34
  From: 10.255.245.43, State: Up, ActiveRoute: 1, LSPname: to-singapore
  ActivePath: long (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Protection: None
Autobandwidth
  MinBW: 104.858Mbps MaxBW: 136.315Mbps AdjustTimer: 300 secs
  Max AvgBW util: 104bps, Bandwidth Adjustment in 215 second(s).
 *Primary   long             State: Up, COS: 6
    Bandwidth: 104.858Mbps
    OptimizeTimer: 250
    Reoptimization in 237 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 50)
          10.35.38.2 S 192.168.135.29 S 10.35.39.1 S 10.35.40.2 S 10.35.41.1 S
    Received RRO:
          10.35.38.2 (flag=0x09) 192.168.135.29 (flag=0x09) 10.35.39.1
   (flag=0x01) 10.35.40.2 (flag=0x01) 10.35.41.1 (flag=0x01)
   25 Dec 11 18:14:34  Change in active path
   24 Dec 11 18:14:34  Record Route:  10.35.38.2 (flag=0x09) 192.168.135.29
   (flag=0x09) 10.35.39.1 (flag=0x09) 10.35.40.2 (flag=0x01) 10.35.41.1
   (flag=0x01)
   23 Dec 11 18:14:34  Up
   22 Dec 11 18:14:34  CSPF: computation result accepted
   21 Dec 11 18:12:10  CSPF: computation result ignored[3 times]
   20 Dec 11 18:01:00  Selected as active path
...
Egress LSP: 1 sessions

10.255.245.43
  From: 10.255.245.34, LSPstate: Up, ActiveRoute: 0, LSPname: to-paris
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  133,  Since: Tue Dec 11 18:06:18 2001
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 17 protocol 0
  PATH rcvfrom: 192.168.135.34  (so-0/0/0.0) 15 pkts
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.35.41.1  10.35.40.2  10.35.39.1  192.168.135.29
192.168.135.34  <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions

10.255.245.36
  From: 10.255.245.39, LSPstate: Up, ActiveRoute: 0, LSPname: skye-path
  Resv style: 1 FF, Label in: 100001, Label out: 3
  Time left:  138,  Since: Tue Dec 11 18:09:46 2001
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 58 protocol 0
  PATH rcvfrom: 192.168.135.13  (so-1/3/1.0) 10 pkts
  PATH sentto: 192.168.135.34 (so-0/0/0.0) 10 pkts
  RESV rcvfrom: 192.168.135.34  (so-0/0/0.0) 10 pkts
  Explct route: 192.168.135.34
  Record route: 192.168.135.13  <self>  192.168.135.34
```

```
10.255.245.40
  From: 10.255.245.39, LSPstate: Up, ActiveRoute: 0, LSPname: cree-path
  Resv style: 1 FF, Label in: 100000, Label out: 100057
  Time left:  128,  Since: Tue Dec 11 18:09:43 2001
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 4 receiver 54 protocol 0
  PATH rcvfrom: 192.168.135.13  (so-1/3/1.0) 11 pkts
  PATH sentto: 10.35.38.2 (so-0/1/0.0) 11 pkts
  RESV rcvfrom: 10.35.38.2  (so-0/1/0.0) 10 pkts
  Explct route: 10.35.38.2 192.168.135.29 10.35.39.1 10.35.40.2
  Record route: 192.168.135.13  <self>  10.35.38.2  192.168.135.29
10.35.39.1  10.35.40.2
Total 2 displayed, Up 2, Down 0
```

show mpls path

Syntax show mpls path *<path-name>*

Description Display the named paths used in dynamic MPLS that have been configured on this router.

Options none—Display all configured LSPs.

path-name—(Optional) Display the specified LSP.

Required Privilege Level view

Output Fields Path name—Information about ingress LSPs. Each path has one line of output.

Address—Addresses of the routers that form the LSP.

Strict/loose address—Whether the address is a configured as a strict or loose address.

Sample Output

```
user@host> show mpls path
Path name        Address          Strict/loose address
p1               123.456.55.6     Strict
                 123.456.1.6      Loose
p2               191.456.1.4      Strict
```

show ted database

Syntax show ted database <brief | detail | extensive> <*system-name*>

Description Display the entries in the Traffic Engineering Database (TED).

Options brief—(Optional) Display brief TED information.

detail—(Optional) Display detailed TED information.

extensive—(Optional) Display all information about the TED.

system-name—(Optional) Display TED information for a particular system.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ted database brief on page 712
Sample Output: show ted database detail on page 712
Sample Output: show ted database extensive on page 712

Options at a Glance Table 89 summarizes which information is included in each show ted database command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 89: Show TED Database Output Fields

Options	Field Description
All	Age (s)—How long since the node was last refreshed, in seconds.
Extensive	Available BW—Amount of bandwidth actually reserved by RSVP for each priority level.
Extensive	Color—Administrative group assigned to this link.
Brief Detail	ID—Host name and address of the node that the link is coming from.
All	LnkIn—Number of nodes pointing towards this node.
All	LnkOut—Number of nodes that this node points towards.
Detail Extensive	Local—Address of the local interface being used to reach remote node.
Extensive	Metric—Configured traffic engineering metric. The default setting is 1.
Extensive	NodeID—Host name and address of the node.
All	Protocol—Protocol that reported the node information.
Detail Extensive	Remote—Address of the interface on the remote node.
Extensive	Reservable bandwidth—Subscription factor for the interface.
Extensive	Static BW—Total interface bandwidth in bps.
All	TED database—Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.
Detail Extensive	To—Address on the far end of a link.
All	Type—Type of node. It can be either Rtr (router) or Net (pseudonode).

Output Fields

TED database—Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.

ID—Host name and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.

ID—(Brief and Detail output only) Host name and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

NodeID—(Extensive output only) Host name and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

Type—Type of node. It can be either Rtr (router) or Net (pseudonode).

Age(s)—How long since the node was last refreshed, in seconds.

LnkIn—Number of nodes pointing towards this node.

LnkOut—Number of nodes to which this node points.

Protocol—Protocol that reported the node information. It can be one or more of the following:

- IS-IS(1)—IS-IS Level 1.
- IS-IS(2)—IS-IS Level 2.
- OSPF (area-*number*)—OSPF from the specified area.

To—(Detail and Extensive output only) Address on the far end of a link.

Local—(Detail and Extensive output only) Address of the local interface being used to reach remote node.

Remote—(Detail and Extensive output only) Address of the interface on the remote node.

Color—(Extensive output only) Administrative group assigned to this link.

Metric—(Extensive output only) Configured traffic engineering metric. The default setting is 1.

Static BW—(Extensive output only) Total interface bandwidth in bps.

Reservable bandwidth—(Extensive output only) Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the subscription statement when configuring RSVP.

Available BW—(Extensive output only) Amount of bandwidth actually reserved by RSVP for each priority level. You can configure this by including the priority statement when configuring LSPs. The bandwidth shown is for the entire interface, not for each individual LSP.

Sample Output: show ted database brief

```
user@host> show ted database brief
TED database: 6 ISIS nodes 6 INET nodes
ID                            Type Age(s) LnkIn LnkOut Protocol
cheviot.00(123.456.1.10)      Rtr    383      1      1 IS-IS(2) IS-IS(1)
corriedale.00(123.456.1.11)   Rtr     36      2      0 IS-IS(2) IS-IS(1)
wolff.00(123.456.1.12)        Rtr    399      0      0 IS-IS(2) IS-IS(1)
perendale.00(123.456.1.13)    Rtr    385      2      0 IS-IS(2) IS-IS(1)
merino.00(123.456.1.14)       Rtr    379      1      3 IS-IS(2) IS-IS(1)
romney.00(123.456.1.15)       Rtr    427      0      2 IS-IS(2) IS-IS(1)
```

Sample Output: show ted database detail

```
user@host> show ted database detail merino
TED database: 6 ISIS nodes 6 INET nodes
NodeID: merino.00(123.456.1.14)
  Type: Rtr, Age: 507 secs, LinkIn: 1, LinkOut: 3
  Protocol: IS-IS(2)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207
    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
  Protocol: IS-IS(1)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207
    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
```

Sample Output: show ted database extensive

```
user@host> show ted database extensive

TED database: 4 ISIS nodes 4 INET nodes
NodeID: strawberry.00(10.255.245.87)
  Type: Rtr, Age: 247 secs, LinkIn: 0, LinkOut: 0
  Protocol: IS-IS(2)
NodeID: orange.00(10.1.1.77)
  Type: Rtr, Age: 46 secs, LinkIn: 1, LinkOut: 1
  Protocol: IS-IS(2)
    To: banana.00(10.255.245.34), Local: 10.0.0.2, Remote: 10.0.0.1
      Color: 0 <none>
      Metric: 10
      Static BW: 622Mbps
      Reservable BW: 622Mbps
      Available BW [priority] bps:
        [0] 622Mbps     [1] 622Mbps     [2] 622Mbps     [3] 622Mbps
        [4] 622Mbps     [5] 622Mbps     [6] 622Mbps     [7] 622Mbps
NodeID:  apple.00(123.456.4.21)
  Type: Rtr, Age: 235 secs, LinkIn: 1, LinkOut: 1
  Protocol: IS-IS(2)
    To: banana.00(10.255.245.34), Local: 123.456.8.209, Remote: 123.456.8.208
      Color: 0 <none>
      Metric: 10
      Static BW: 155Mbps
      Reservable BW: 155Mbps
      Available BW [priority] bps:
        [0] 155Mbps     [1] 155Mbps     [2] 155Mbps     [3] 155Mbps
        [4] 155Mbps     [5] 155Mbps     [6] 155Mbps     [7] 155Mbps
  Protocol: IS-IS(1)
```

```
NodeID: banana.00(10.255.245.34)
  Type: Rtr, Age: 247 secs, LinkIn: 2, LinkOut: 2
  Protocol: IS-IS(2)
    To:  apple.00(123.456.4.21), Local: 123.456.8.208, Remote: 123.456.8.209
      Color: 0 <none>
      Metric: 10
      Static BW: 155Mbps
      Reservable BW: 155Mbps
      Available BW [priority] bps:
        [0] 155Mbps     [1] 155Mbps     [2] 155Mbps     [3] 155Mbps
        [4] 155Mbps     [5] 155Mbps     [6] 155Mbps     [7] 155Mbps
    To: orange.00(10.1.1.77), Local: 10.0.0.1, Remote: 10.0.0.2
      Color: 0 <none>
      Metric: 10
      Static BW: 622Mbps
      Reservable BW: 622Mbps
      Available BW [priority] bps:
        [0] 622Mbps     [1] 622Mbps     [2] 622Mbps     [3] 622Mbps
        [4] 622Mbps     [5] 622Mbps     [6] 622Mbps     [7] 622Mbps
```

show ted link

Syntax show ted link <brief | detail>

Description Display traffic engineering database (TED) link information.

Options brief—(Optional) Display brief TED link information.

detail—(Optional) Display detailed TED link information.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ted link brief on page 715
Sample Output: show ted link detail on page 715

Options at a Glance Table 90 summarizes which information is included in each show ted link command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 90: Show TED Link Output Field Summary

Options	Field Description
Detail	–>*host-name*—Host name and address of the node that the link is going to.
Brief	–>ID—Host name and address of the node that the link is going to.
Detail	*host-name*—Host name and address of the node that the link is coming from.
Brief	ID—Host name and address of the node that the link is coming from.
All	Local BW—How much bandwidth the local router has placed on the link.
All	Local Path—How many paths CSPF on the local router has placed on the link.

Output Fields ID—(Brief output only) Host name and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

–>ID—(Brief output only) Host name and address of the node that the link is going to. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

host-name—(Detail output only) Host name and address of the node that the link is coming from. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

–>*host-name*—(Detail output only) Host name and address of the node that the link is going to. An address of .00 indicates that the node is the router itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.

Local Path—How many paths CSPF on the local router has placed on the link.

Local BW—How much bandwidth the local router has placed on the link.

Sample Output: show ted link brief

```
user@host> show ted link brief
TED link:
ID                          ->ID                          LocalPath LocalBW
cheviot.00(123.456.1.10)     merino.00(123.456.1.14)             0 0bps
merino.00(123.456.1.14)      corriedale.00(123.456.1.11)         0 0bps
merino.00(123.456.1.14)      perendale.00(123.456.1.13)          0 0bps
merino.00(123.456.1.14)      cheviot.00(123.456.1.10)            0 0bps
romney.00(123.456.1.15)      corriedale.00(123.456.1.11)         0 0bps
romney.00(123.456.1.15)      perendale.00(123.456.1.13)          0 0bps
```

Sample Output: show ted link detail

```
user@host> show ted link detail
TED link:
cheviot.00(123.456.1.10)->merino.00(123.456.1.14), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->cheviot.00(123.456.1.10), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
romney.00(123.456.1.15)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
romney.00(123.456.1.15)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
```

show ted protocol

Syntax show ted protocol <brief | detail>

Description Display information about the protocols from which the traffic engineering database (TED) learned about its nodes.

Options brief—(Optional) Display brief protocol information.

detail—(Optional) Display detailed protocol information.

Default: brief

Required Privilege Level view

Output Fields Protocol name—Protocol that reported the node information. It can be one of the following:

- IS-IS(1)—IS-IS Level 1.
- IS-IS(2)—IS-IS Level 2.
- OSPF (area-*number*)—OSPF from the specified area.

Credibility—If there is conflicting information about a node, rank the protocols by credibility. The protocol with the highest credibility value is the one that the TED uses.

Self node—Address the protocol uses as the local address.

Sample Output

```
user@host> show ted protocol
Protocol name        Credibility  Self node
IS-IS(2)             2  (highest) corriedale.00(123.456.1.11)
IS-IS(1)             1            corriedale.00(123.456.1.11)
```

Chapter 43
RSVP Monitoring and Troubleshooting

Table 91 summarizes the command-line interface (CLI) commands you can use to monitor the RSVP. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 91: Commands for Monitoring RSVP

Task or Information to Monitor	Command
Status of interfaces on which RSVP is running	show rsvp interface on page 719
Neighbors with which this router exchanges information	show rsvp neighbor on page 724
Currently active RSVP sessions	show rsvp session on page 727
Clear all RSVP sessions	clear rsvp session on page 717
RSVP packet and error counters	show rsvp statistics on page 734
Clear RSVP packet and error counters	clear rsvp statistics on page 718
RSVP version and configuration information	show rsvp version on page 736

clear rsvp session

Syntax clear rsvp session <name *name*>

Description Tear down RSVP sessions. RSVP PathTear and ResvTear messages tear down the sessions.

Options none—Reset and restart all RSVP sessions that this router knows about; that is, all sessions for which this router is the ingress, transit, or egress router. Depending on the number of sessions involved, it might take a while to tear down all the RSVP sessions.

name *name*—(Optional) Reset and restart the specified RSVP session.

Required Privilege Level clear

See Also clear mpls lsp on page 692

Sample Output

```
user@host> clear rsvp session
user@host>
```

clear rsvp statistics

Syntax clear rsvp statistics

Description Clear RSVP packet and error counters.

Required Privilege Level clear

show rsvp interface

Syntax show rsvp interface <brief | detail | link-management >

Description Display the status of each interface on which RSVP is enabled, and display packet statistics for each interface.

Options brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

link-management—(Optional) Display the control peers and corresponding TE-link information created by Link Management Protocol (LMP).

Default: brief

Required Privilege Level view

Sample Output Sample Output: show rsvp interface brief on page 722
Sample Output: show rsvp interface detail on page 722
Sample Output: show rsvp interface (standard, link-management) on page 722

Options at a Glance Table 92 summarizes the information included in the output of each show rsvp interface command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 92: Show RSVP Interface Output Fields Summary

Options	Field Description
Detail	Ack—Acknowledge message for refresh reductions.
Standard	Active control channel—Next-hop TE-link address to transmit messages.
All	Active resv—Number of reservations that are actively reserving bandwidth on the interface.
Detail	Address—IP address of the local interface.
All	Available BW—Amount of bandwidth that RSVP is allowed to reserved, in bps.
Detail	End-to-End RSVP—Statistics for the number of end-to-end RSVP messages sent.
Detai	Hello—Number of RSVP hello packets that have been sent to and received from the neighbor.
Detail	HelloInt—Frequency at which RSVP hellos are sent on this interface (in seconds). HelloInt can have a value from 0-60 seconds. The default is 3 seconds.
All	Highwater mark—Highest bandwidth that has ever been reserved on this interface, in bps.
Detail	Index—Index of the interface.
All	Interface—Name of the interface.
Detail	Last 5 seconds Received—Number of packets received in the last 5 seconds.
Detail	Last 5 seconds Sent—Number of packets sent in the last 5 seconds.
Standard	Local identifier—Identifier of the local side of the link.
Detail	NoAggregate—There is no refresh reduction support on the interface.
Detail	NoAuthentication—The interface does not support RSVP authentication.
Detail	NoLinkProtection—There is no link protection on the interface.
Detail	NoReliable—There is no refresh reduction message-id extension supported on the interface.
Detail	PacketType—Type of RSVP packet.
Detail	Path—Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.

Options	Field Description
Detail	PathErr—Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
Detail	PathTear—Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.
Detail	PreEmptionCnt—Number of times an RSVP session was presented on this interface.
All	Reserved BW—Currently reserved bandwidth in bps.
Detail	Resv—Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and which store reservation state information in each node along the path.
All	RSVP interface—Number of interfaces on which RSVP is active.
Detail	Srefresh—Summary refresh messages.
All	State—State of the interface. It can be Enabled, Disabled, Up, or Down.
All	Static BW—Total interface bandwidth in bps.
All	Subscription—User-configured subscription factor. The default is 100%.
Standard	TElink—Traffic-engineered links that are managed by the peer they are associated with.
Detail	Total Received—Total number of packets received since RSVP was enabled.
Detail	Total Sent—Total number of packets sent.
Detail	Update threshold—Percentage change in reserved bandwidth to trigger an IGP update.

Output Fields

RSVP interface—Number of interfaces on which RSVP is active. Each interface has one line of output.

Interface—Name of the interface.

Index—Index of the interface.

State—State of the interface.

- Disabled—No TE info is displayed.
- Down—The interface is not operational.
- Enabled—Displays TE info.
- Up—The interface is operational.

NoAuthentication—The interface does not support RSVP authentication.

NoAggregate—There is no refresh reduction support on the interface.

NoReliable—There is no refresh reduction message-id extension supported on the interface.

NoLinkProtection—There is no link protection on the interface.

HelloInterval—Frequency at which RSVP hellos are sent on this interface (in seconds). HelloInt can have a value from 0-60 seconds. The default is 3 seconds.

Address—IP address of the local interface.

Active control channel—Nexthop TE-link address to transmit messages.

TElink—Traffic-engineered links that are managed by the peer they are associated with.

Local identifier— Identifier of the local side of the link.

Active resv—Number of reservations that are actively reserving bandwidth on the interface.

PreEmptionCnt—Number of times an RSVP session was presented on this interface.

Update threshold—Percentage change in reserved bandwidth to trigger an IGP update. The configured value can be 1-20%. The default is 10%.

Subscription—User-configured subscription factor. The default is 100%.

Static BW—Total interface bandwidth in bps.

Available BW—Amount of bandwidth that RSVP is allowed to reserved, in bps. It is equal to (static bandwidth * subscription factor).

Reserved BW—Currently reserved bandwidth in bps.

Highwater mark—Highest bandwidth that has ever been reserved on this interface, in bps.

PacketType—(Detail output only) Type of RSVP packet.

Total Sent—(Detail output only) Total number of packets sent.

Total Received—(Detail output only) Total number of packets received since RSVP was enabled.

Last 5 seconds Sent—(Detail output only) Number of packets send in the last 5 seconds.

Last 5 seconds Received—(Detail output only) Number of packets received in the last 5 seconds.

Path—Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.

PathErr—Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.

PathTear—Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.

Resv—Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and which store reservation state information in each node along the path.

End-to-End RSVP—Statistics for the number of end-to-end RSVP messages sent.

Hello—Number of RSVP hello packets that have been sent to and received from the neighbor.

Ack—Acknowledge message for refresh reductions.

Srefresh—Summary refresh messages.

Sample Output: show rsvp interface brief

```
user@host> show rsvp interface brief
RSVP interface: 1 active
                  Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
de0.0     Up         1    23%  10Mbps      989.992kbps 1.31Mbps    1.31Mbps
```

Sample Output: show rsvp interface (standard, link-management)

```
user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C  State: Up
Active Control Channel: so-0/1/0.0

  TElink: TElnk1, Link ID: 37811
  ActiveResv 0, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

  TElink: TElnk2, Link ID: 37808
  ActiveResv 1, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B  State: Up
Active Control Channel: so-1/0/0.0

  TElink: TElnkAB1, Link ID: 1598
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

  TElink: TElnkAB2, Link ID: 1597
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

Sample Output: show rsvp interface detail

```
user@host> show rsvp interface detail
so-0/1/1.0  Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  PacketType              Total                  Last 5 seconds
                   Sent      Received        Sent      Received
  Path               16             0           1             0
  PathErr             0             0           0             0
  PathTear            1             0           0             0
  Resv                0            11           0             1
  ResvErr             0             0           0             0
  ResvTear            0             0           0             0
  Hello              66            67           1             1
  Ack                 0             0           0             0
  Srefresh            0             0           0             0
  EndtoEnd RSVP       0             0           0             0
```

```
fe-0/0/0.0  Index 4, State: Dis/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.129
  PacketType              Total                  Last 5 seconds
                   Sent      Received        Sent      Received
  Path                 0             0           0             0
  PathErr              0             0           0             0
  PathTear             0             0           0             0
  Resv                 0             0           0             0
  ResvErr              0             0           0             0
  ResvTear             0             0           0             0
  Hello                0             0           0             0
  Ack                  0             0           0             0
  Srefresh             0             0           0             0
  EndtoEnd RSVP        0             0           0             0
```

show rsvp neighbor

Syntax show rsvp neighbor <detail>

Description Display a list of RSVP neighbors that were learned dynamically when exchanging RSVP packets. Once a neighbor is learned, it is never removed from the list of RSVP neighbors.

none—Display RSVP neighbor information.

detail—Display detailed RSVP neighbor information.

Required Privilege Level view

Output Fields RSVP neighbor—Number of neighbors about which the router has learned. Each neighbor has one line of output.

via—(Detail only) Name of the interface where the neighbor has been detected.

Address—Address of a learned neighbor.

Idle—Amount of time the neighbor has been idle, in seconds.

Up/Dn—Number of neighbor up/down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running JUNOS software Release 3.2 or earlier, are not reported as up or down.

Up cnt and Down cnt—(Detail output only) Number of neighbor up/down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running JUNOS software Release 3.2 or earlier, are not reported as up or down.

status—(Detail output only) State of the physical interface.

- Up—Router can detect RSVP Hello messages from the neighbor.
- Down—Router cannot detect RSVP Hello messages from the neighbor.

LastChange—Time elapsed since the neighbor state changed either from up to down or from down to up. The format is hh:mm:ss.

Last changed time—(Detail output only) Time elapsed since the neighbor state changed either from up to down or from down to up. The format is hh:mm:ss.

HelloInt—Frequency at which RSVP hellos are sent on this interface (in seconds). HelloInt can have a value from 0-60 seconds. The default is 3 seconds.

HelloTx/Rx—Number of hello packets that have been sent to and received from the neighbor.

Hello—(Detail output only) Number of RSVP hello packets that have been sent to and received from the neighbor.

MsgRcvd—Number of Path and Resv messages that this router has received from the neighbor.

Message received—(Detail output only) Number of Path and Resv messages that this router has received from the neighbor.

Remote Instance—(Detail output only) Identification provided by the remote router during Hello message exchange.

Local Instance—(Detail output only) Identification sent to the remote router during Hello message exchange.

Refresh reduction—(Detail output only) Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve router performance by reducing the process overhead, thus increasing the number of LSPs a router can support. Refresh reduction can have the following values:

- operational—All 4 RSVP refresh reduction extensions (RFC 2961 message ack, bundling, summary refresh and staged refresh timer) are functional between the two neighboring routers.
- incomplete—Some of the RSVP refresh reduction extensions are functional between the two neighboring routers.
- no operational—Either the refresh reduction feature has been turned off, or the remote router cannot support the refresh reduction extensions.

Remote end—(Detail output only) Neighboring router's status in regard to refresh reduction. It can have the following values:

- enabled—Remote router has requested refresh reduction during RSVP message exchanges.
- disabled—Remote router does not require refresh reduction.

Ack-extension—(Detail output only) An RSVP refresh reduction extension. It can have the following values:

- enabled—Both local and remote routers support the ack-extension (RFC 2961).
- disabled—Remote router does not support the ack-extension.

Link protection—(Detail output only) Status of the MPLS fast reroute mechanism that protects traffic from link failure. It can have the following values:

- enabled—Link protection feature has been turned on, protecting the neighbor with a bypass LSP.
- disabled—No link protection feature has been enabled for this neighbor.

LSP name—(Detail output only) Name of the bypass LSP.

Bypass LSP—(Detail output only) Status of the bypass LSP. It can have the following values:

- does not exist—Bypass LSP is not available.
- connecting—The router is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment.
- operational—Bypass LSP is up and running.
- down—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.

Backup routes—(Detail output only) Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).

Backup LSPs—(Detail output only) Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).

Bypass explicit route—(Detail output only) The Explicit Route Object's (ERO) path that is taken by the bypass LSP.

Sample Output: show rsvp neighbor (standard)

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address            Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203       0  3/2       13:01        3   366/349
192.168.207.207       0  1/0       22:49        3   448/448
```

Sample Output: show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: so-0/1/0.0    status: Up
   Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
   Message received: 632
   Hello: sent 673, received 656, interval 3 sec
   Remote instance: 0x6432838a, Local instance: 0x74b72e36
   Refresh reduction:  operational
     Remote end: enabled,  Ack-extension: enabled
   Link protection:  enabled
     LSP name: Bypass_to_192.168.207.203
     Bypass LSP: operational,  Backup routes: 1,  Backup LSPs: 0
     Bypass explicit route: 192.168.207.207 192.168.207.224
```

show rsvp session

Syntax show rsvp session <terse | brief | detail> <bidirectional | unidirectional> <interface *interface*> <*lsp-type*> <name *session-name*> <*session-type*> <te-link *te-link*> <up | down>

Description Display information about RSVP sessions.

Options brief—(Optional) Display brief session information.

detail—(Optional) Display detailed session information.

bidirectional—(Optional) Display only bidirectional LSPs. The default is to display both unidirectional and bidirectional LSPs.

down—(Optional) Display only LSPs that are inactive.

interface *interface*—(Optional) Display sessions with reservations on the specified interface.

lsp-type—(Optional) Display information about a particular LSP. Because RSVP sessions optionally can be used for MPLS label-switched-path negotiations, this option filters out those sessions that are used for setting up LSPs.

- lsp—Sessions used to set up LSPs.
- nolsp—Sessions not used to set up LSPs.

name *session-name*—(Optional) Display information about the named session.

session-type—(Optional) Display information about a particular session type:

- egress—Sessions that terminate on this router.
- ingress—Sessions that originate from this router.
- transit—Sessions that transit through this router.

te-link *te-link*—(Optional) Display sessions with reservations on the specified TE-link.

terse—(Optional) Display only a count of the number of sessions, with no information about individual sessions. Specifying this option is useful when there is a large number of sessions.

unidirectional—(Optional) Displays only unidirectional LSPs. The default is to display both unidirectional and bidirectional LSPs.

up—(Optional) Display only LSPs that are active.

Default: brief; display active and inactive LSPs

Required Privilege Level view

Sample Output

Sample Output: show rsvp session terse on page 731
Sample Output: show rsvp session (brief, standard) on page 731
Sample Output: show rsvp session detail on page 731

Options at a Glance

Table 93 summarizes the information included in the output of each show rsvp session command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 93: Show RSVP Session Output Field Summary

Options	Field Description
Detail	...incomplete—The recorded route information is incomplete, possibly because the session is not established.
Detail	Active Route—Number of active routes (prefixes) that have been installed in the forwarding table.
Detail	address—(First line of each group) Destination (egress router) of the LSP.
Brief Standard	Bidir—The LSP allows data to travel in both directions between GMPLS devices.
Detail	Bidirectional—The LSP allows data to travel both ways between GMPLS devices.
Brief Terse Detail	Egress RSVP—Information about egress RSVP sessions.
Detail	Fast reroute desired—Either Frame Relay or Link Protection has been requested by the ingress router.
Brief	From—Source (ingress router) of the session.
Detail	From—Source (ingress router) of the session.
Brief Terse Detail	Ingress RSVP—Information about ingress RSVP sessions.
Brief Detail	Label in—Incoming label for this LSP.
Brief Detail	Label out—Outgoing label for this LSP.
Brief Detail	LSP name—Name of the LSP.
Brief Detail	LSP state—State of the LSP that is being handled by this RSVP session. It can be either Up or Dn (down).
Detail	PATH rcvfrom—Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor.
Detail	PATH sentto—Address of the next-hop (downstream) router or client, interface used to reach this neighbor, and how many packets were sent to the downstream router.
Detail	Port number—Protocol ID and sender/receiver port used in this RSVP session.
Detail	Record route—Recorded route for the session as taken from the Record Route Object.
Detail	RESV rcvfrom—Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor. The output in this field, which is consistent with that in the PATH field, indicates that the RSVP negotiation is complete.
Brief Detail	Resv style—RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).
Brief	Rt—Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).
Detail	Since—Date and time when the RSVP session was initiated.
Detail	Suggested label received—Label the upstream node suggests to use in the Resv message that is sent. Displayed only when the LSP is bidirectional.
Detail	Suggested label sent—Label the downstream node suggests to use in its Resv message that is returned. Displayed only when the LSP is bidirectional.

Options	Field Description
Brief Detail	Time left—Number of seconds remaining in the lifetime of the reservation.
Brief	To—Destination (egress router) of the session.
Brief Terse Detail	Transit RSVP—Information about transit RSVP sessions.
Detail	Tspec—Sender's traffic specification, which describes the sender's traffic parameters.
Detail	Type—The LSP type can be Bypass LSP, Link protected LSP, Backup LSP at Point-of-Local-Repair, or Backup LSP at Merge Point.
Detail	Upstream label in—Incoming label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.
Detail	Upstream label out—Outgoing label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Output Fields

Ingress RSVP—(Detail output only) Information about ingress RSVP sessions.

Ingress RSVP—(Brief and Terse output only) Information about ingress RSVP sessions. Each session has one line of output.

Egress RSVP—Information about egress RSVP sessions.

Transit RSVP— Information about the transit RSVP sessions.

To—(Brief output only) Destination (egress router) of the session.

From—(Brief output only) Source (ingress router) of the session.

address—(Detail output only; first line of each group) Destination (egress router) of the LSP.

From—(Detail output only) Source (ingress router) of the session.

LSP state—(Brief and detail output only) State of the LSP that is being handled by this RSVP session. It can be either Up or Dn (down).

Rt—(Brief output only) Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).

Active Route—(Detail output only) Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).

LSP name—(Brief and detail output only) Name of the LSP.

Bidir—The LSP will allow data to travel in both directions between GMPLS devices.

Bidirectional—The LSP will allow data to travel both ways between GMPLS devices.

Upstream label in—Incoming label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Upstream label out—Outgoing label for reverse direction traffic for this LSP. Displayed only when the LSP is bidirectional.

Suggested label received —Label the upstream node suggests to use in the Resv message that is sent. Displayed only when the LSP is bidirectional.

Suggested label sent—Label the downstream node suggests to use in its Resv message that is returned. Displayed only when the LSP is bidirectional.

Resv style—(Brief and detail output only) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).

Label in—(Brief and detail output only) Incoming label for this LSP.

Label out—(Brief and detail output only) Outgoing label for this LSP.

Time left—(Brief and detail output only) Number of seconds remaining in the lifetime of the reservation.

Since—(Detail output only) Date and time when the RSVP session was initiated.

Tspec—(Detail output only) Sender's traffic specification, which describes the sender's traffic parameters.

Port number—(Detail output only) Protocol ID and sender/receiver port used in this RSVP session.

Fast reroute desired—(Detail output only) Either Frame Relay or Link Protection has been requested by the ingress router.

Type—(Detail output only) The LSP type. It can be:

- Bypass LSP—LSP that is used to protected one or more user LSPs in case of link failure.
- Link protected LSP—LSP has been protected by a bypass LSP at the outgoing interface.
- Backup LSP at Point-of-Local-Repair (PLR)—LSP that has been temporarily established to protected a user LSP at the ingress of a failed link.
- Backup LSP at Merge Point (MP)—LSP that has been temporarily established to protected a user LSP detected at the egress of a failed link.

PATH rcvfrom—(Detail output only) Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor.

PATH sentto—(Detail output only) Address of the next-hop (downstream) router or client, interface used to reach this neighbor, and how many packets were sent to the downstream router.

RESV rcvfrom—(Detail output only) Address of the previous-hop (upstream) router or client, interface the neighbor used to reach this router, and how many packets were received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.

Explct route—(Detail output only) The Explicit Route Object's (ERO) path that is taken by the bypass LSP.

Record route—(Detail output only) Recorded route for the session as taken from the Record Route Object.

Sample Output: show rsvp session terse

```
user@host> show rsvp session terse
Ingress RSVP: 2 sessions
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
Total 1 displayed, Up 1, Down 0
```

Sample Output: show rsvp session (brief, standard)

```
user@host> show rsvp session
Ingress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.196  10.255.245.194  Up     1  1 SE       -   100001 pro3-ac
10.255.245.196  10.255.245.194  Up     0  1 FF       -    22276 Glsp1 Bidir
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

Sample Output: show rsvp session detail

```
user@host> show rsvp session detail
Ingress RSVP: 2 sessions

10.255.245.196
  From: 10.255.245.194, LSPstate: Up, ActiveRoute: 1
  LSPname: pro3-ac
  Resv style: 1 SE, Label in: -, Label out: 100001
  Time left:    -,  Since: Mon Apr  8 12:12:56 2002
  Tspec: rate 20Mbps size 20Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 16 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 192.168.207.30 (so-0/1/1.0) 14 pkts
  RESV rcvfrom: 192.168.207.30  (so-0/1/1.0) 12 pkts
  Explct route: 192.168.207.30 192.168.207.38
  Record route: <self>  192.168.207.30  192.168.207.38
```

```
10.255.245.196
  From: 10.255.245.194, LSPstate: Up, ActiveRoute: 0
  LSPname: Glsp1
  Bidirectional, Upstream label in: 22278, Upstream label out: -
  Suggested label received: -, Suggested label sent: 22276
  Resv style: 1 FF, Label in: -, Label out: 22276
  Time left:    -,  Since: Mon Apr  8 12:11:28 2002
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 1 receiver 19 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 10.255.245.196 (so-0/1/1.0) 17 pkts
  RESV rcvfrom: 192.168.207.38  (PEER-C) 12 pkts
  Explct route: 2.2.2.2
  Record route: <self>  2.2.2.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 2 sessions

10.255.245.194
  From: 10.255.245.195, LSPstate: Up, ActiveRoute: 0
  LSPname: Gpro3-ba
  Bidirectional, Upstream label in: -, Upstream label out: 39811
  Suggested label received: 39811, Suggested label sent: -
  Resv style: 1 FF, Label in: 39811, Label out: -
  Time left:  154,  Since: Mon Apr  8 12:11:38 2002
  Tspec: rate 0bps size 0bps peak 622.08Mbps m 20 M 1500
  Port number: sender 1 receiver 17 protocol 0
  PATH rcvfrom: 192.168.207.132  (fe-0/0/0.0) 13 pkts
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 4.4.4.2  <self>

10.255.245.194
  From: 10.255.245.195, LSPstate: Up, ActiveRoute: 0
  LSPname: pro3-ba
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left:  149,  Since: Mon Apr  8 12:12:43 2002
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 19 protocol 0
  PATH rcvfrom: 192.168.207.26  (so-1/0/0.0) 13 pkts
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 192.168.207.26  <self>
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions

10.255.245.198
  From: 10.255.245.197, LSPstate: Up, ActiveRoute: 0
  LSPname: pro3-de
  Resv style: 1 SE, Label in: 100000, Label out: 3
  Time left:  117,  Since: Mon Apr  8 12:12:56 2002
  Tspec: rate 50Mbps size 50Mbps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 16 protocol 0
  FastReroute desired
  PATH rcvfrom: 192.168.207.130  (fe-0/0/0.0) 11 pkts
  PATH sentto: 192.168.207.30 (so-0/1/1.0) 13 pkts
  RESV rcvfrom: 192.168.207.30  (so-0/1/1.0) 12 pkts
```

```
  Explct route: 192.168.207.30
  Record route: 192.168.207.130  <self>  192.168.207.30
    Detour is Up
    Detour PATH sentto: 192.168.207.34 (so-0/1/0.0) 12 pkts
    Detour RESV rcvfrom: 192.168.207.34  (so-0/1/0.0) 11 pkts
    Detour Explct route: 192.168.207.34
    Detour Record route: 192.168.207.130  <self>  192.168.207.34
Total 1 displayed, Up 1, Down 0
```

show rsvp statistics

Syntax show rsvp statistics

Description Display RSVP packet and error statistics.

Required Privilege Level view

Output Fields Packet Type—Statistics about different RSVP messages.

Total Sent—Total number of packets sent since RSVP was enabled.

Total Received—Total number of packets received since RSVP was enabled.

Last 5 seconds Sent—Total number of packets sent in the last 5 seconds.

Last 5 seconds Received—Number of packets received in the last 5 seconds.

Path—Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.

PathErr—Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.

PathTear—Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.

Resv FF—Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.

Resv WF—Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.

Resv SE—Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.

ResvErr—Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.

ResvTear—Statistics about ResvTear messages, which remove reservation states along a path.

ResvConf—Statistics about ResvConfirm messages, which are responses to confirm a reservation request.

Ack—Acknowledge message for refresh reductions.

SRefresh—Summary refresh messages.

Hello—Number of RSVP hello packets that have been sent to and received from the neighbor.

EndtoEnd RSVP—Statistics for the number of End-to-end RSVP messages.

Errors—Statistics about errored RSVP packets.

Rcv pkt bad length—The packet was not processed because its length is inappropriate.

Rcv pkt unknown type—The packet is not one of the well-known RSVP types, as defined in RFC 2205.

Rcv pkt bad version—The packet is not an RSVP version 1 packet.

Rcv pkt auth fail—The packet failed authentication checks.

Rcv pkt bad cksum—The RSVP checksum check failed.

Rcv pkt bad format—General packet processing failed because the packet was badly formed.

Memory alloc fail—An internal resource failure occurred.

No path info—A reservation was received, but no sender is active.

Resv style conflict—The same session contains inconsistent reservation styles.

Port conflict—There were inconsistent port numbers for the same session.

Resv no interface—An interface for the receive reservation packets cannot be located.

PathErr to client—Number of PathErr packets delivered to the local client.

ResvErr to client—Number of ResvErr packets delivered to the local client.

Path timeout—Number of times the sender timed out because the path was removed.

Resv timeout—Number of times the receiver timed out because the reservation was removed.

Message out-of-order—Records the number of RSVP incoming messages that are considered out-of-order. Message out-of-order is detected from the message-id object's sequence number.

Unknown ack msg—A neighboring router replies an ACK object that contains an unknown message-id. This can indicate a message-id handshake problem.

Recv nack—A neighboring router explicitly rejects a message ID in a summary refresh message. This can happen if that neighbor has been rebooted. In this case, the router sends a regular RSVP refresh message to recover the state, and starts the message-id handshake process again.

Recv duplicated msg-id—Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring router restarts.

Sample Output

```
user@host> show rsvp statistics
   PacketType              Total                  Last 5 seconds
                     Sent      Received        Sent      Received
   Path               355           408           0             0
   PathErr              2            13           0             0
   PathTear           101           139           0             0
   Resv FF              0             0           0             0
   Resv WF              0             0           0             0
   Resv SE            419           225           0             0
   ResvErr              0             0           0             0
   ResvTear             0            13           0             0
   ResvConf             0             0           0             0
   Ack                682          1414           0             0
```

```
SRefresh        395198        236030           5             2
Hello           578809        578221           4             4
EndtoEnd RSVP        0             0           0             0

Errors                          Total            Last 5 seconds
Rcv pkt bad length                0                          0
Rcv pkt unknown type              0                          0
Rcv pkt bad version               0                          0
Rcv pkt auth fail                 0                          0
Rcv pkt bad checksum              0                          0
Rcv pkt bad format                0                          0
Memory allocation fail            0                          0
No path information               10                          0
Resv style conflict               0                          0
Port conflict                     0                          0
Resv no interface                 0                          0
PathErr to client                38                          0
ResvErr to client                 0                          0
Path timeout                      8                          0
Resv timeout                     57                          0
Message out-of-order              0                          0
Unknown ack msg                2978                          0
Recv nack                        86                          0
Recv duplicated msg-id            5                          0
```

show rsvp version

Syntax show rsvp version

Description Display information about the version of the RSVP software and the settings of the refresh timer and keep multiplier.

Required Privilege Level view

Sample Output

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
   RSVP protocol       = Enabled
   R (refresh timer)   = 30 seconds
   K (keep multiplier) = 3
```

Chapter 44
LDP Monitoring and Troubleshooting

Table 94 summarizes the command-line interface (CLI) commands you can use to monitor the LDP. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 94: Commands for Monitoring LDP

Task or Information to Monitor	Command
Entries in the LDP database	show ldp database on page 738
Status of interfaces on which LDP is running	show ldp interface on page 741
LDP neighbor	show ldp neighbor on page 742
Clear LDP neighbors	clear ldp neighbor on page 737
Configured named paths that are used by LDP	show ldp path on page 743
LDP routing table entries	show ldp route on page 745
Currently active LDP sessions	show ldp session on page 747
LDP statistics	show ldp statistics on page 749
Clear LDP statistics	clear ldp statistics on page 737

clear ldp neighbor

Syntax clear ldp neighbor <*address*>

Description Tear down neighbor connections.

Options none—Tear down connections with all LDP neighbors.

address—(Optional) Tear down the connection with the specified neighbor.

Required Privilege Level clear

clear ldp statistics

Syntax clear ldp statistics

Description Zero all LDP statistics.

Required Privilege Level clear

show ldp database

Syntax show ldp database <brief | detail | extensive> < instance *instance-name* >

Description Display entries in the LDP label database.

Options brief—(Optional) Display brief information about entries in the LDP database.

detail—(Optional) Display detailed information about the entries in the LDP database.

extensive—(Optional) Display extensive information about the entries in the LDP database. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp database brief on page 740
Sample Output: show ldp database detail on page 740

Options at a Glance Table 95 summarizes which information is included in the output of each show ldp database command option. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 95: Show LDP Database Output Fields Summary

Options	Field Description
All	Input label database—Label received from the other router.
All	Label—Label binding to a route prefix.
All	Output label database—Label advertised to the other router.
All	Prefix—Route prefix.
All	*session-identifier*—Session identifier, which comprises the local and remote label space identifiers.
Detail	State—State of the label binding. It can be Active, New, MapRcv, MapSend, RelRcv, RelRsnd, RelSend, ReqSend, or W/dSend

Output Fields Input label database—Label received from the other router.

Output label database—Label advertised to the other router.

session-identifier—Session identifier, which comprises the local and remote label space identifiers.

Label—Label binding to a route prefix.

Prefix—Route prefix. It can be either the IP prefix or the Layer 2 encapsulation type in the format L2CKT *encapsulation type* VC *number*, where *encapsulation type* can be:

- FRAME RELAY
- ATM AAL5
- ATM CELL
- VLAN
- ETHERNET
- CISCO_HDLC
- PPP

The virtual circuit *number* can have any numeric value.

State—(Detail output only) State of the label binding. It can be:

- Active—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.
- New—New label that has not yet been distributed.
- MapRcv—Waiting to receive a label mapping message.
- MapSend—Waiting to send a label mapping message.
- RelRcv—Waiting to receive a label release message.
- RelRsnd—Waiting to receive a label release message before resending label mapping message.
- RelSend—Waiting to send a label release message.
- ReqSend—Waiting to send a label request message.
- W/dSend—Waiting to send a label withdrawal message.

Sample Output: show ldp database brief

```
user@host> show ldp database brief
Input label database, 10.255.245.222:0--10.255.245.221:0
  Label     Prefix
 100018     10.255.245.222/32
      3     10.255.245.221/32
 100011     L2CKT FRAME RELAY VC 11

Output label database, 10.255.245.222:0--10.255.245.221:0
  Label     Prefix
 100018     10.255.245.222/32
      3     10.255.245.221/32
 100011     L2CKT FRAME RELAY VC 1
```

Sample Output: show ldp database detail

```
user@host> show ldp database detail
Input label database, 10.255.245.222:0--10.255.245.221:0
  Label     Prefix
 100018     10.255.245.222/32
            State: Active
      3     10.255.245.221/32
            State: Active
 100011     L2CKT FRAME RELAY VC 1
            State: Active

Output label database, 10.255.245.222:0--10.255.245.221:0
  Label     Prefix
      3     10.255.245.222/32
            State: Active
 100019     10.255.245.221/32
            State: Active
 100011     L2CKT FRAME RELAY VC 1
            State: Active
```

show ldp interface

Syntax show interface <brief | detail | extensive> < instance *instance-name* >

Description Display the status of each interface on which LDP is enabled.

Options none—Display all interfaces in the LDP database.

brief—(Optional) Display brief interface information.

detail—(Optional) Display detailed interface information.

extensive—(Optional) Display extensive interface information. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp interface brief on page 741
Sample Output: show ldp interface detail on page 741

Output Fields Interface—Interface name.

Label space ID—Label space identifier that the router is advertising on the interface.

Nbr count—Number of neighbors on the interface.

Next hello—How long until the next hello packet is sent on this interface, in seconds.

Hello interval—(Detail output only) Configured hello interval, in seconds.

Hold time—(Detail output only) Configured hold time, in seconds.

Sample Output: show ldp interface brief

```
user@host> show ldp interface brief
Interface            Label space ID        Nbr count   Next hello
fxp0.0               10.10.255.6:0            2           3
mps0.0               10.10.255.6:0            0           0
```

Sample Output: show ldp interface detail

```
user@host> show ldp interface detail
Interface            Label space ID        Nbr count   Next hello
fxp0.0               10.10.255.6:0            2           1
  Hello interval: 5, Hold time: 15
mps0.0               10.10.255.6:0            0           3
  Hello interval: 5, Hold time: 15
```

show ldp neighbor

Syntax show ldp neighbor <brief | detail | extensive> < instance *instance-name* >

Description Display a list of LDP neighbors.

Options brief—(Optional) Display brief neighbor information.

detail—(Optional) Display detailed neighbor information.

extensive—(Optional) Display extensive neighbor information. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp neighbor brief on page 742
Sample Output: show ldp neighbor detail on page 742

Output Fields Address—IP address of the neighbor.

Interface—Interface over which the neighbor was discovered.

Label space ID—Label space identifier advertised by the neighbor.

Hold time—Remaining hold time before the neighbor expires, in seconds.

Transport address—(Detail output only) Address to which the neighbor wants the local route to establish the LDP session.

Configuration sequence—(Detail output only) Counter that increments whenever the neighbor changes its configuration.

Sample Output: show ldp neighbor brief

```
user@host> show ldp neighbor brief
Address            Interface          Label space ID         Hold Time
192.168.1.213      so-0/0/0           10.10.255.4:0            13
192.168.1.211      so-0/0/0           10.10.255.2:0            14
```

Sample Output: show ldp neighbor detail

```
user@host> show ldp neighbor detail
Address            Interface          Label space ID         Hold Time
192.168.1.213      so-0/0/0           10.10.255.4:0            10
  Transport address: 10.10.255.4, Configuration sequence: 3
192.168.1.211      s0-0/0/0           10.10.255.2:0            11
  Transport address: 10.10.255.2, Configuration sequence: 1
```

show ldp path

Syntax show ldp path <brief | detail | extensive> < instance *instance-name* >

Description Display the LSPs that LDP created.

Options none— Display all LSPs.

brief—(Optional) Display brief information about the LSPs on this router.

detail—(Optional) Display detailed information about the LSPs on this router.

extensive—(Optional) Display extensive information about the LSPs on this router. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp path brief on page 743
Sample Output: show ldp path detail on page 744

Output Fields Output Session (label)—Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.

Input Session (label)—Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.

route—Each line in the output corresponds to an MPLS route. For an example, see "Sample Output: show ldp path detail". The fourth line shows that any MPLS packet received from 10.10.255.6 with label 100000 will be transmitted to 10.10.255.3 with label 3.

Attached route—(Detail output only) Route corresponding to the LSP.

Ingress route—(Detail output only) Indicates that the router acts as the ingress for the LSP.

Sample Output: show ldp path brief

```
user@host> show ldp path brief
Output Session (label)          Input Session (label)
10.10.255.6:0(3)                (egress)
10.10.255.3:0(3)                   .
10.10.255.4:0(3)                   .
10.10.255.6:0(100000)           10.10.255.3:0(3)
10.10.255.4:0(100000)              .
10.10.255.3:0(100001)           10.10.255.6:0(3)
10.10.255.4:0(100001)              .
10.10.255.6:0(100002)           10.10.255.4:0(3)
10.10.255.3:0(100002)              .
```

Sample Output: show ldp path detail

```
user@host> show ldp path detail
Output Session (label)          Input Session (label)
10.10.255.6:0(3)                (egress)
10.10.255.3:0(3)                   .
10.10.255.4:0(3)                   .
  Attached route: *10.10.255.2/32
10.10.255.6:0(100000)           10.10.255.3:0(3)
10.10.255.4:0(100000)              .
  Attached route: *10.10.255.3/32, Ingress route
10.10.255.3:0(100001)           10.10.255.6:0(3)
10.10.255.4:0(100001)              .
  Attached route: *10.10.255.6/32, Ingress route
10.10.255.6:0(100002)           10.10.255.4:0(3)
10.10.255.3:0(100002)              .
  Attached route: *10.10.255.4/32, Ingress route
```

show ldp route

Syntax show ldp route <brief | detail | extensive> < instance *instance-name* >

Description Display the entries in the LDP routing table. These are the unicast routes that LDP is considering for binding to LSPs. This routing table is a subset of the unicast routing table, inet.0.

Options brief—(Optional) Display brief information about the entries in the LDP routing table.

detail—(Optional) Display detailed information about the entries in the LDP routing table.

extensive—(Optional) Display extensive information about the entries in the LDP routing table. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp route brief on page 745
Sample Output: show ldp route detail on page 746

Output Fields Destination—Destination prefix.

Next-hop intf—Interface that is the next hop to the destination prefix.

Next-hop address—IP address of the next hop.

Bound to outgoing label—(Detail output only) Indicates the route has been bound to LSPs with the label being distributed for that LSP.

Sample Output: show ldp route brief

```
user@host>show ldp route brief
Destination              Next-hop intf    Next-hop address
 10.10.255.1/32          so-2/3/0
*10.10.255.3/32          so-1/0/0         10.10.1.3
*10.10.255.1/32          so-2/3/0
 10.10.255.4/32          so-0/0/0         192.168.1.213
*10.10.255.4/32          so-0/0/0         192.168.1.213
 10.10.255.6/32          so-0/0/0         192.168.1.215
*10.10.255.6/32          so-0/0/0         192.168.1.215
*10.10.255.2/32
 0.0.0.0/0               so-0/0/0         192.168.1.254
 10.10.255.3/32          so-1/0/0         10.10.1.3
```

Sample Output: show ldp route detail

```
user@host> show ldp route detail
Destination              Next-hop intf    Next-hop address
 10.10.255.1/32          so-2/3/0
*10.10.255.3/32          so-1/0/0         10.10.1.3
   Bound to outgoing label 100001
*10.10.255.1/32          so-2/3/0
 10.10.255.4/32          so-0/0/0         192.168.1.213
*10.10.255.4/32          so-0/0/0         192.168.1.213
   Bound to outgoing label 100002
 10.10.255.6/32          so-0/0/0         192.168.1.215
*10.10.255.6/32          so-0/0/0         192.168.1.215
   Bound to outgoing label 100000
*10.10.255.2/32
   Bound to outgoing label 3
 0.0.0.0/0               so-0/0/0         192.168.1.254
 10.10.255.3/32          so-1/0/0         10.10.1.3
```

show ldp session

Syntax show ldp session <brief | detail | extensive> < instance *instance-name* >

Description Display information about LDP sessions.

Options none—Display all LSP sessions.

brief—(Optional) Display brief LDP session information.

detail—(Optional) Display detailed LDP session information.

extensive—(Optional) Display extensive LDP session information. The output of the extensive option contains information for Juniper Networks Customer Support personnel only.

instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ldp session brief on page 748
Sample Output: show ldp session detail on page 748

Output Fields Address—Transport address of the session.

State—State of the session. It can be Nonexistent, Connecting, Initialized, OpenRec, OpenSent, Operational, or Closing, which correspond to the state diagram specified in the LDP functional specification Internet draft.

Connection—TCP connection state. It can be Closed, Opening, or Open.

Hold time—Remaining hold time before the session will be closed, in seconds.

Session ID—(Detail output only) Label space identifier of the session.

Next keepalive—(Detail output only) Time until next keepalive is sent, in seconds.

Active—(Detail output only) Whether the local router is playing the active role in the session.

Maximum PDU—(Detail output only) Maximum PDU size for the session.

Hold time—(Detail output only) Remaining hold time before the session will be closed, in seconds.

Neighbor count—(Detail output only) Number of neighbors that are contributing to the session.

Keepalive interval—(Detail output only) Keepalive interval, in seconds.

Connect retry interval—(Detail output only) TCP connection retry interval, in seconds.

Local address—(Detail output only) Local transport address.

Remote address—(Detail output only) Remote transport address.

Next-hop addresses received—(Detail output only) Next-hop addresses received on the session.

Sample Output: show ldp session brief

```
user@host> show ldp session brief
  Address           State        Connection     Hold Time
10.10.255.2         Operational  Open            25
10.10.255.4         Operational  Open            26
```

Sample Output: show ldp session detail

```
user@host> show ldp session detail

Address: 10.10.255.2, State: Operational, Connection: Open, Hold time: 23
  Session ID: 10.10.255.6:0--10.10.255.2:0
  Next keepalive in 4 seconds
  Active, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Keepalive interval: 10, Connect retry interval: 10
  Local address: 10.10.255.6, Remote address: 10.10.255.2
  Next-hop addresses received:
    10.10.1.2
    10.10.255.2
    192.168.1.211

Address: 10.10.255.4, State: Operational, Connection: Open, Hold time: 24
  Session ID: 10.10.255.6:0--10.10.255.4:0
  Next keepalive in 9 seconds
  Active, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Keepalive interval: 10, Connect retry interval: 10
  Local address: 10.10.255.6, Remote address: 10.10.255.4
  Next-hop addresses received:
    10.10.1.4
    10.10.255.4
    192.168.1.213
```

show ldp statistics

Syntax show ldp statistics < instance *instance-name* >

Options: instance *instance-name*—(Optional) Display routing instance information for the specified instance. If *instance-name* is omitted, then the master instance is assumed.

Description Display LDP statistics.

Required Privilege Level view

Output Fields Message type—LDP message types.

Total Sent, Received—Total number of each message type sent and received.

Last 5 seconds Sent, Received—Number of each message type send and received in the last 5 seconds.

Event type—LDP events and errors.

Total—Total number of each event or error.

Last 5 seconds—Number of each event or error in the last 5 seconds.

Sample Output

```
user@host> show ldp statistics

Message type              Total                     Last 5 seconds
                       Sent      Received          Sent      Received
Hello                   265           263             2             2
Initialization            2             2             0             0
Keepalive               112           111             1             0
Notification              0             0             0             0
Address                   2             2             0             0
Address withdraw          0             0             0             0
Label mapping             7             6             0             0
Label request             0             0             0             0
Label withdraw            2             0             0             0
Label release             0             2             0             0
Label abort               0             0             0             0
All UDP                 265           263             2             2
All TCP                 123           121             1             0

Event type                            Total        Last 5 seconds

Sessions opened                           2                 0
Sessions closed                           0                 0
Topology changes                         11                 0

No interface                              0                 0
No session                                0                 0
No adjacency                              0                 0
Unknown version                           0                 0
Malformed PDU                             0                 0
Malformed message                         0                 0
Unknown message type                      0                 0
Inappropriate message                     0                 0
Malformed TLV                             0                 0
Bad TLV value                             0                 0
Missing TLV                               0                 0
PDU too large                             0                 0
```

Chapter 45

Layer 2 VPN Monitoring and Troubleshooting

Table 96 summarizes the command-line interface (CLI) commands you can use to monitor Layer 2 VPNs. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

These commands have been condensed and grouped to specifically address VPNs, with references to the complete command elsewhere in this book.

Table 96: Commands for Monitoring Layer 2 VPNs

Task Category	Task or Information to Monitor	Command
BGP Information	Entries in the BGP neighbor database	show bgp neighbor on page 752
	BGP summary information.	show bgp summary on page 757
Interface Information	General interface information.	show interfaces terse on page 759
Layer 2 Connections	Display Layer 2 virtual circuit information.	show l2circuit connections on page 760
	Display Layer 2 VPN information.	show l2vpn connections on page 765
Routing Table Information	Information about the entries in the routing tables.	show route on page 768
	Routes transmitted by a particular routing protocol.	show route advertising-protocol on page 769
	Information about the entries in the kernel's forwarding table.	show route forwarding-table on page 771
	Routes received by a particular routing protocol.	show route receive-protocol on page 774
	Routes in a particular routing table.	show route table on page 777

show bgp neighbor

Syntax show bgp neighbor <*neighbor-address*>

Description Display information about BGP peers.

Options none—Display information about all BGP peers.

neighbor-address—(Optional) IP address of a BGP peer.

Required Privilege Level view

See Also For information about the local-address, nlri, hold-time, and preference statements, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols.*

Output Fields Peer—Address of the BGP neighbor. The address is followed by the neighbor's port number.

AS—AS number of the peer.

Local—Address of the local router. The address is followed by the peer's port number.

Type—Type of peer. It can be either Internal or External.

State—Current state of the BGP session. It can be one of the following:

- Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message.
- Connect—BGP is waiting for the transport protocol connection to complete.
- Established—The BGP session has been established, and the peers are exchanging update messages.
- Idle—This is the first stage of a connection. BGP is waiting for a Start event.
- OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.
- OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Flags—Internal BGP flags. It can be one or more of the following:

- CleanUp—The peer session is being shut down.
- Delete—This peer has been deleted.
- Idled—This peer has been permanently idled.
- Initializing—The peer session is initializing.
- SendRtn—Messages are being sent to the peer.
- TryConnect—Another attempt is being made to connect to the peer.

- Unconfigured—This peer is not configured.
- WriteFailed—An attempt to write to this peer failed.

Last State—Previous state of the BGP session.

Last Event—Last activity that occurred in the BGP session. It can be one of the following:

- Closed—The BGP session closed.
- ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect.
- HoldTime—The session ended because the hold timer expired.
- KeepAlive—The local router sent a BGP keepalive message to the peer.
- Open—The local router sent a BGP open message to the peer.
- OpenFail—The local router did not receive an acknowledgment of a BGP open message from the peer.
- RecvKeepAlive—The local router received a BGP keepalive message from the peer.
- RecvNotify—The local router received a BGP notification message from the peer.
- RecvOpen—The local router received a BGP open message from the peer.
- RecvUpdate—The local router received a BGP update message from the peer.
- Start—The peering session started.
- Stop—The peering session stopped.
- TransportError—A TCP error occurred.

Last Error—Last error that occurred in the BGP session. It can be one of the following:

- Cease—An error occurred, such as a version mismatch, that caused the session to close.
- Finite State Machine Error—In setting up the session, BGP received a message that it did not understand.
- Hold Time Expired—The session's hold time expired.
- Message Header Error—The header of a BGP message was malformed.
- Open Message Error—A BGP open message contained an error.
- None—No errors occured in the BGP session.
- Update Message Error—A BGP update message contained an error.

Export—Name of the export policy that is configured on the peer.

Import—Name of the import policy that is configured on the peer.

Options—Configured BGP options. It can be one or more of the following:

- AddressFamily—Configured address family. It can be inet or inet-vpn.
- GracefulRestart—Graceful restart is configured.
- HoldTime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.
- Local Address—Address configured with the local-address statement.
- NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any.
- Peer AS—Configured peer AS.
- Preference—Preference value configured with the preference statement. The default preference value is 170.
- Refresh—Configured to refresh automatically when policy changes.
- Rib-group—Configured routing table group.

Address families configured—Names of configured address families for VPN.

Local Address—The interface used by an external BGP session (except when you use the BGP multihop option for this peer).

Holdtime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.

Preference—Preference value configured with the preference statement. The default preference value is 170.

Number of flaps—Number of times the BGP session has gone down and then come back up.

Peer ID—Router identifier of the peer.

Local ID—Router identifier of the local router.

Active hold time—Hold time the local router negotiated with the peer.

Keepalive interval—Keepalive interval, in seconds.

Local Interface—Name of directly connected interface.

NLRI for restart configured on peer—Names of address families configured for restart.

NLRI advertised by peer—Address families supported by the peer. It can be unicast or multicast.

NLRI for this session—Address families being used for this session.

Peer supports Refresh capability—Remote peer supports Route Refresh Capability, and is able to send and request full route table readvertisements. For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.

Restart time configured on the peer—Configured time allowed for restart on the neighbor. The range is 1-600 seconds. The default value is 300 seconds.

Stale routes from peer are kept for—When graceful restart is negotiated, this is the maximum time allowed to hold routes from neighbors after the BGP session has gone down. The range is 1-600 seconds. The default value is 300 seconds.

Restart time requested by this peer—Restart time requested by this neighbor during capability negotiation.

NLRI that peer supports restart for—Address families indicates graceful restart during capability negotiation.

NLRI that restart is negotiated for—Both address families are supported by both this peer and its neighbor.

NLRI of received end-of-rib markers—Address families for which end-of-rib markers are received from the neighbor.

NLRI of all end-of-rib markers sent—Address families for which end-of-rib markers are sent to the neighbor.

Table inet.*number*—Information about the routing table.

- RIB State—BGP is in the graceful restart process for this routing table. It can be restart is complete or restart in progress.
- Bit—Number that represents the entry in the routing table for this peer.
- Send state—State of the BGP group. It can be in sync, not in sync, or not advertising.
- Active prefixes—Number of prefixes received from the peer that are active in the routing table.
- Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table.
- Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

Last traffic (seconds)—Last time any traffic was received from the peer or sent to the peer, and the last time the local router checked.

Input messages—Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output messages—Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output queue[0], Output queue[1]—Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.

Trace options—Configured tracing of BGP protocol packets and operations. See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Trace file—Name of the file to receive the output of the tracing operation.

Sample Output: show bgp neighbor (PE side)

```
user@host> show bgp neighbor bgp.l2.vpn.0
Peer: 192.168.16.1+1037 AS 65412 Local: 192.168.24.1+179 AS 65412
  Type: Internal    State: Established    Flags: <>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast l2vpn
  Local Address: 192.168.24.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 192.168.16.1     Local ID: 192.168.24.1     Active Holdtime: 90
  Keepalive Interval: 30
  NLRI advertised by peer: inet-unicast inet-multicast l2vpn
  NLRI for this session: inet-unicast l2vpn
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 0
    Suppressed due to damping: 0
  Table bgp.l2vpn.0 Bit: 30000
    Send state: in sync
    Active prefixes: 1
    Received prefixes: 1
    Suppressed due to damping: 0
  Table vpna.l2vpn.0 Bit: 50000
    Send state: in sync
    Active prefixes: 1
    Received prefixes: 1
```

show bgp summary

Syntax show bgp summary

Description Display summary information about BGP and its neighbors. Show if routes were received from all of the remote PE routers.

Required Privilege Level view

Output Fields Groups—Number of BGP groups.

Peers—Number of BGP peers.

Down peers—Number of down BGP peers.

Table—Name of routing table.

Tot Paths—Total number of paths.

Act Paths—Number of active routes.

Suppressed—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

History—Number of withdrawn routes stored locally to keep track of damping history.

Damp State—Number of routes with a Figure of Merit (FOM) greater than 0, but still active because the FOM has not reached the threshold at which suppression occurs.

Pending—Routes in process by BGP import policy.

Peer—Address of each BGP peer. Each peer has one line of output.

AS—Peer's AS number.

InPkt—Number of packets received from the peer.

OutPkt—Number of packets sent to the peer.

OutQ—Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.

Flaps—Number of times the BGP session has gone down and then come back up.

Last Up/Down—Time since the neighbor transitioned to or from the established state.

State|#Active/Received/Damped—Displays either the BGP state or, if the neighbor is in the established state, the number of paths received from the neighbor, as well as the number of these paths that have been accepted as active and hence are being used for forwarding.

Sample Output: show bgp summary

```
user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State
Pending
bgp.l2vpn.0            1          1          0          0
0          0
inet.0                 0          0          0          0
0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35   65299         72         74       0       1       19:00
Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299       2164       2423       0       4       19:50
Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37   65299         36         37       0       4       17:07
Establ
  inet.0: 0/0/0
10.255.245.39   65299        138        168       0       6       53:48
Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.69   65299        134        140       0       6       53:42
Establ
  inet.0: 0/0/0
```

show interfaces terse

Syntax show interfaces terse <*interface-name*>

Description Display summary information about interfaces.

Options *interface-name*—(Optional) Name of an interface.

Required Privilege Level view

Output Fields Interface—Name of the interface.

Admin—Whether the interface is turned on (up) or off (down).

Link—Link state. It can be Up or Down.

Proto—Protocol configured on that interface.

Local—Local address of the interface.

Remote—For point-to-point interfaces, address of the remote side of the connection.

Sample Output: show interfaces terse

```
user@host> show interfaces terse |match ccc
so-0/3/0.0      up    up   ccc
```

Sample Output: show interfaces terse

```
user@host> run show interfaces terse |match so
so-0/0/0        up    up
so-0/0/0.1      up    down inet  10.1.13.2/30
                           iso
so-0/0/0.2      up    down inet  10.1.23.2/30
                           iso
so-0/0/0.4      up    down inet  10.1.34.1/30
                           iso
so-0/0/0.5      up    up   inet  10.1.35.1/30
                           iso
so-0/0/1        up    up
so-0/0/2        up    up
so-0/0/3        up    up
                           iso
                           iso
                           iso
                           iso
                           iso
                           iso
47.0005.80ff.f800.0000.0108.0001.0102.5524.5219.00
```

show l2circuit connections

Syntax show l2circuit connections <neighbor ip-address><*status*><brief | extensive>

Description Display status information about Layer 2 virtual circuits from the local PE router to its neighbors.

Options ping—Display information about all Layer 2 virtual circuits to all neighbors. The standard output is identical to the output when the status option is used.

brief—(Optional) Display brief Layer 2 virtual circuit information.

extensive—(Optional) Display very detailed Layer 2 virtual circuit information.

neighbor ip-address—(Optional) Display all virtual circuits between the PE router and the specified neighbor.

status—(Optional) Display information about the connection and interface status. It can be one of the following:

- up—Display information about all VCs that are in the up state.
- down—Display information about all VCs that are in the down state.
- up-down—Display information about all VCs that are in both the up and down states.

Required Privilege Level view

Sample Output Sample Output: show l2circuit connections brief on page 763
Sample Output: show l2circuit connections extensive on page 763
Sample Output: show l2circuit connections status on page 764

Options at a Glance Table 97 summarizes which information is included in the output of each of the show l2circuit connections command options. In this table, output fields are listed in alphabetical order. In the Output Fields section, the output fields are listed in the order in which they are displayed.

Table 97: Show l2circuit connections Output Field Summary

Options	Field Description
Brief Extensive	#Uptrans—Number of times the VC came up.
Brief Extensive	(vc *number*)—virtual circuit identifier. A unique number between two PE routers that identifies the VC configured between them.
Extensive	event—Event types logged in history.
Extensive	Incoming label—Label used by the remote side of the VC to send packets destined to the local side.
Brief Extensive	Interface—Logical PE-to-CE interface on which the virtual circuit is configured.
Extensive	Interface/L—Name associated with an event.
Brief Extensive	Layer-2 Circuit Connections—Displays the legends for connection and interface status.
Extensive	Local interface—Name of the local PE-to-CE interface.
Brief Extensive	Neighbor—Remote PE neighbor.
Extensive	Outgoing label—Label used by the local side of the VC to send packets to the remote side of the VC.
Extensive	Remote PE—Prefix of the remote PE router.
Brief Extensive	St—Status of the VC connection.
Extensive	Status—Status of the local interface.
Extensive	time—Time of the occurrence of the event.
Brief Extensive	Time last up—Date and time the VC was last operational.
Brief Extensive	Type—VC type. It can be rmt (remote) or loc (local).

Output Fields

Layer-2 Circuit Connections—Displays the legends for connection and interface status.

Neighbor—Remote PE neighbor.

Interface—Logical PE-to-CE interface on which the virtual circuit is configured.

Type—The VC can be rmt (remote) or loc (local).

St—Status of the VC connection:

- EI—The local VC interface is configured with an encapsulation which is not supported.
- EM—Encapsulation type received on this VC from the neighbor does not match the local VC interface encapsulation type.
- OL—No advertisement has been received for this VC from the neighbor. There is no outgoing label available for use by this VC.
- Dn—VC is down because the local VC interface is down.
- VC-Dn—Virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor.
- UP—Operational.
- XX—VC is down for an unknown reason. This is a programming error.

Time last up—Date and time the VC was last operational.

#Uptrans—Number of times the VC came up.

Local interface—Name of the local PE-to-CE interface.

Status—Status of the local interface.

- UP—Interface is operational.
- Dn—Interface is not operational.
- NP—(not present) The interface doesn't exist.
- DS—Disabled. Interface has been administratively disabled.
- WE—Wrong encapsulation. The interface is not configured as CCC.
- UN—Interface status is initialized.

Remote PE—Prefix of the remote PE router.

Incoming label—Label used by the remote side of the VC to send packets destined to the local side. This label is routed to the local VC interface.

Outgoing label—Label used by the local side of the VC to send packets to the remote side of the VC. Packets originated on the local VC interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this VC. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this VC.

time—Time of the occurrence of the event.

event—Event types logged in history.

- loc intf up—Local VC interface went up.
- loc intf down—Local VC interface went down.
- In lbl Update—Incoming label is updated.
- Out lbl Update—Outgoing label is updated.
- PE route changed—Route to PE router is updated.
- PE route down—Route to PE router is down.
- rmt side marked—Remote side is marked.
- VC Dn—Remote side indicated that its end of the VC is down (if the tunnel LSP from the remote side to the local side is down).
- status update timer—Status update timer processing. It computes the state of the VC, and determines whether it should be advertised or withdrawn to or from the remote side.

Interface/L—Name associated with an event. It can be the name of the interface, the name of the new label, or if the route to a PE router changed, the name of the PE router that went down.

(vc *number*)—Virtual circuit identifier is a unique number (between 1 and 2^{32}-1) between 2 PE routers that identifies the VC configured between them. Between 2 PE routers, the virtual circuit identifier is unique, but not across the entire network. A virtual circuit identifier in conjunction with a neighbor address uniquely identifies the VC.

Sample Output: show l2circuit connections brief

```
user@host> show l2circuit connections brief
Layer-2 Circuit Connections:

Legend for connection status (St)   Legend for interface status
EI -- encapsulation invalid         UP -- operational
EM -- encapsulation mismatch        Dn -- down
OL -- no outgoing label             NP -- no present
Dn -- down                          DS -- disabled
VC-Dn -- Virtual circuit Down       WE -- wrong encapsulation
UP -- operational                   UN -- uninitialized
XX -- unknown

Neighbor: 10.1.1.195
    Interface                 Type  St     Time last up          # Up trans
    so-0/1/0.1 (vc 1)         rmt   Up     Oct  8 15:44:19 2001           1
    so-0/1/0.2 (vc 2)         rmt   Up     Oct  8 15:44:19 2001           1
```

Sample Output: show l2circuit connections extensive

```
user@host> show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)   Legend for interface status
EI -- encapsulation invalid         UP -- operational
EM -- encapsulation mismatch        Dn -- down
OL -- no outgoing label             NP -- no present
Dn -- down                          DS -- disabled
VC-Dn -- Virtual circuit Down       WE -- wrong encapsulation
UP -- operational                   UN -- uninitialized
XX -- unknown

Neighbor: 10.1.1.195
    Interface                 Type  St     Time last up          # Up trans
    so-0/1/0.1 (vc 1)         rmt   Up     Oct  8 15:44:19 2001           1
      Local interface: so-0/1/0.1, Status: Up
      Remote PE: 10.1.1.195
      Incoming label: 100002, Outgoing label: 138264604
        Time                  Event                  Interface/Lbl/PE
        Oct  8 15:44:19 2001  status update timer
        Oct  8 15:44:18 2001  PE route changed
        Oct  8 15:44:18 2001  Out lbl Update           138264604
        Oct  8 15:44:18 2001  In lbl Update               100002
        Oct  8 15:44:18 2001  loc intf up             so-0/1/0.1
    so-0/1/0.2 (vc 2)         rmt   Up     Oct  8 15:44:19 2001           1
      Local interface: so-0/1/0.2, Status: Up
      Remote PE: 10.1.1.195
      Incoming label: 100003, Outgoing label: 138264604
        Time                  Event                  Interface/Lbl/PE
        Oct  8 15:44:19 2001  status update timer
        Oct  8 15:44:17 2001  PE route changed
        Oct  8 15:44:17 2001  Out lbl Update           138264604
        Oct  8 15:44:17 2001  In lbl Update               100003
        Oct  8 15:44:17 2001  loc intf up             so-0/1/0.2
```

Sample Output: show l2circuit connections status

```
user@host> show l2circuit connections status
Layer-2 Circuit Connections:

Legend for connection status (St)   Legend for interface status
EI -- encapsulation invalid         UP -- operational
EM -- encapsulation mismatch        Dn -- down
OL -- no outgoing label             NP -- no present
Dn -- down                          DS -- disabled
VC-Dn -- Virtual circuit Down       WE -- wrong encapsulation
UP -- operational                   UN -- uninitialized
XX -- unknown

Neighbor: 10.1.1.195
    Interface                 Type  St     Time last up          # Up trans
    so-0/1/0.1 (vc 1)         rmt   Up     Oct  8 15:44:19 2001           1
      Local interface: so-0/1/0.1, Status: Up
      Remote PE: 10.1.1.195
      Incoming label: 100002, Outgoing label: 138264604
    so-0/1/0.2 (vc 2)         rmt   Up     Oct  8 15:44:19 2001           1
      Local interface: so-0/1/0.2, Status: Up
      Remote PE: 10.1.1.195
      Incoming label: 100003, Outgoing label: 138264604
```

show l2vpn connections

Syntax show l2vpn connections <down | up | up-down> <history | status | brief | extensive>
<instance *instance* | local-site *local-site* | remote-site *remote-site*>

Description Display Layer 2 VPN connections.

Options brief—(Optional) Display one line of connection status.

down—(Optional) Display Layer 2 VPN connections that are not operational.

extensive—(Optional) Display Layer 2 VPN connection status and history.

history—(Optional) Display Layer 2 VPN connection history.

instance—(Optional) Name of the routing instance.

local-site—(Optional) Display Layer 2 VPN local-site name or ID.

remote-site—(Optional) Display Layer 2 VPN remote-site name or ID.

status—(Default) Display Layer 2 VPN connection and circuit status.

up—(Optional) Display Layer 2 VPN connections that are operational.

up-down—(Optional) Display both operational and non-operational Layer 2 VPN connections.

Required Privilege Level view

Sample Output Sample Output: show l2vpn connections on page 767
Sample Output: show l2vpn connections extensive on page 767

Output Fields Instance—Name of Layer 2 VPN instance.

Local site—Name of local site.

Interface name—Name of interface.

Remote Site ID—Remote site ID.

Label Offset—Label offset.

Offset—Offset.

Range—Range.

connection-site—Name of connection site.

Type—Type of connection. The value can be:

- OR—Out of range up.
- EI—Encapsulation invalid.
- EM—Encapsulation mismatch.
- CN—Circuit not present.
- OL—No outgoing label.
- Dn—Down.
- VC-Dn—Virtual circuit down.
- -> —Only outbound connection is up.
- <- —Only inbound connection is up.
- UP—Operational.
- XX—Unknown.

St—Status of connection.

Time last up—Time connection was last in the Up condition.

Up trans—Number of transitions from Down to Up condition.

Local circuit—Address and status of local circuit.

Remote circuit—Address and status of remote circuit.

Status—Status of Local or remote circuit. The value can be:

- up—Operational
- Dn—Down
- NP—No present
- DS—Disabled
- WE—Wrong encapsulation
- UN—Uninitialized

Time—Date and time of Layer 2 VPN connection event.

Event—Type of event.

Interface/Lbl/PE—Interface, Label, PE.

Sample Output: show l2vpn connections

```
user@host> show l2vpn connections
L2VPN Connections :
Instance : vpna
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
    connection-site           Type  St  Time last up          # Up trans
    3 (3)                     loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
      Remote PE: 192.168.16.1
      Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
    connection-site           Type  St  Time last up          # Up trans
    2 (ce-b)                  loc   Up  Jul 18 20:45:46 2001           1
      Local circuit: fe-0/0/0.1, Status: Up
      Remote circuit: fe-0/0/3.0, Status: Up
    1                         rmt   Up  Jul 18 21:47:25 2001           1
      Local circuit: fe-0/0/3.1, Status: Up
      Remote PE: 192.168.16.1
      Incoming label: 33792, Outgoing label: 32770
```

Sample Output: show l2vpn connections extensive

```
user@host> show l2vpn connections extensive
L2VPN Connections:
Legend for connection status (St)   Legend for circuit status
OR -- out of range                   up -- operational
EM -- encapsulation mismatch        Dn -- down
CN -- circuit not present            NP -- no present
OL -- no outgoing label              DS -- disabled
Dn -- down                           WE -- wrong encapuslation
VC-Dn -- Virtual circuit Down        UN -- uninititalized
-> -- only outbound conn is up
<- -- only inbound  conn is up
UP -- operational
XX -- unknown

Instance: vpn-a
Local site: ce-a (1)
  Interface name      Remote Site ID
  fe-0/0/0.0                       2
  Label Offset      Offset     Range
         32768           1         2
    connection-site           Type  St  Time last up          # Up trans
    2                         rmt   Up  Aug  3 00:08:14 2001           1
      Local circuit: fe-0/0/0.0, Status: Up
      Remote PE: 192.168.24.1
      Incoming label: 32769, Outgoing label: 32768
        Time                  Event            Interface/Lbl/PE
        Aug  3 00:08:14 2001  PE route up
        Aug  3 00:08:14 2001  Out lbl Update               32768
        Aug  3 00:08:14 2001  In lbl Update                32769
        Aug  3 00:08:14 2001  ckt0 up                 fe-0/0/0.0
```

show route

Syntax show route <*destination-prefix*> <all | hidden> <brief | detail | extensive> <*filters*>

Description Display the active entries in the routing tables. Checks if the Layer 2 VPN is up and if the remote CE is connected.

Options none—Display all active entries in the routing tables.

all—(Optional) Display information about all routes, including hidden entries.

brief—(Optional) Display brief route information.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

detail—(Optional) Display detailed route information.

extensive—(Optional) Display very detailed route information.

filters—(Optional) One or more filters that select specific information to display from the routing tables. To simplify the readability and explanation of the filters, each filter is treated as an individual command even though you can combine filters in a single show route command. The following is the syntax of the show route command including all filters:

```
show route <destination-prefix> <summary | detail | extensive> <best | exact | range>
   <advertising-protocol protocol neighbor-address> <aspath-regex regular-expression>
   <community as-number:community-value> <damped> <detail> <forwarding-table>
   <inactive> <next-hop address> <output interface-name> <protocol protocol> <range>
   <receive-protocol protocol neighbor-address> <source-gateway address>
   <table routing-table>
```

hidden—(Optional) Display hidden route information.

Default: brief

Required Privilege Level view

Output Fields See the output fields for show route on page 516.

Sample Output See the sample output for show route on page 516.

show route advertising-protocol

Syntax show route advertising-protocol *protocol neighbor-address* <detail>

Description Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol. The information displayed reflects the routes that the routing table has exported into the routing protocol and that were filtered by that protocol's export routing policy statements. (In the figure "Importing and Exporting Routing Policies" in the section "How to Construct Routing Policy" in the *JUNOS Internet Software Configuration Guide: Policy Framework*, these are the routes shown in the right "Protocol" box.)

Options detail—(Optional) Display additional information that was advertised to the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were advertised, including communities, route reflection, and aggregator.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route. It can be bgp or rip.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Nexthop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

Route-distinguisher—The unique route distinguisher associated with this routing instance.

VPN Label—Name of label assigned to this VPN.

Communities—Community path attribute of the route.

Sample Output

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail

vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
     Route Distinguisher: 192.168.16.1:1
     Label-base : 32768, range : 3
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:65412:100 Layer2-info: encaps:VLAN, control
flags:0, mtu: 0
```

show route forwarding-table

Syntax show route forwarding-table <*destination-prefix*> <family *family* | matching *matching*> <multicast> <vpn *vpn*> <detail | summary>

Description Display the route entries in the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine.

The show route forwarding-table command displays only the network-layer prefixes and their next hops. Commonly, you need information in the show route commands. However, when you run into problems, such as connectivity problems, you might need to use the show route forwarding-table command to verify that the routing protocol process has relayed the correct information into the forwarding table.

> **Note** The show route forwarding-table command is an independent command, not a filter that selects specific information that is displayed from the routing tables. You cannot use this command in conjunction with any of the show route filter options (see Table 65, "Commands for Monitoring Protocol-Independent Routing Properties" on page 512).

Options none—Display the routes in the forwarding table. This option is equivalent to the UNIX netstat –rn command.

destination-prefix—(Optional) IP address of a destination.

detail—(Optional) Display detailed information, including clone routes. This option is equivalent to the UNIX netstat –rna command.

family *family*—(Optional) Display routing table entries for the specified family.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

summary—(Optional) Display summary information about routes.

vpn *vpn*—(Optional) Display routing table entries for specified VPN.

Required Privilege Level view

Output Fields Internet, ISO, MPLS—IP, ISO, and MPLS addresses.

Destination—Destination of the route.

Type—How the route was placed into the forwarding table (located in the second column):

- clon—Clone route (for TCP or multicast only; displayed only if you specify the detail option).
- dest—Remote addresses directly reachable through an interface.
- iddn—dest route for which the interface is down.

- ifcl—clon route for which the interface is down.
- ifdn—intf route for which the interface is down.
- ignr—Ignore this route.
- intf—Installed as the result of configuring an interface.
- perm—Permanent (installed by kernel when the routing table is initialized).
- user—Installed by the routing protocol process or as a result of the configuration.

RtRef—IP, ISO, and MPLS addresses.

InIf—Incoming interface.

Flags—Additional information about a route.

- 0x01—Static route.
- 0x02—Cache route.
- 0x04—Check against incoming interface.
- 0x08—Route has accounting.
- 0x10—Route has been sent to the Packet Forwarding Engine.

Nexthop—Next hop to the destination.

Type—Information about the next hop (located in the fourth column):

- bcst—Broadcast.
- dscd—Discard; no ICMP unreachable message sent.
- hold—Next hop is waiting to be resolved (will turn into a unicast or multicast).
- locl—Local address on an interface.
- mcrt—“Regular” multicast next hop.
- mcst—Wire multicast next hop (limited to the LAN).
- mdsc—Multicast discard.
- mgrp—Multicast group member.
- recv—Receive.
- rjct—Discard; ICMP unreachable message sent.

- rslv—Resolving next hop.
- ucst—Unicast.
- ulst—List of unicast next hops. A packet sent to this next hop will go to any next hop in the list.

Index—Number associated with the next hop.

NhRef—Number of routes that reference this next hop.

Netif—Interface used to reach the next hop.

indr—An index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.

Sample Output: show route forwarding-table (| find)

```
user@host>show route forwarding-table | find ccc
Routing table:: ccc
MPLS:
Interface.Label    Type RtRef Nexthop       Type Index NhRef Netif
default            perm     0               dscd     3     1
0                  user     0               recv     5     2
1                  user     0               recv     5     2
32769              user     0               ucst    45     1 fe-0/0/0.534
fe-0/0/0. (CCC)    user     0               indr    44     2
                              10.0.16.2     Push 32768, Push
100004(top)fe-0/0/1.0
```

show route receive-protocol

Syntax show route receive-protocol *protocol neighbor-address* <detail | extensive >

Description Display the routing information as it was received through a particular neighbor of a particular dynamic routing protocol for each VPN. This information includes the routes that the local router advertised to the neighbor. The information displayed reflects the routes before they are filtered by that protocol's import routing policy statements and is placed into the routing table.

Options detail—(Optional) Display additional information that was received from the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were received, including communities, route reflection, and aggregator.

extensive—(Optional) Display extensive information that was received from the neighbor.

neighbor-address—Address of the neighboring router from which the route entry was received.

protocol—Protocol transmitting the route. It can be either bgp or rip.

Required Privilege Level view

Sample Output Sample Output: show route receive-protocol bgp detail on page 775
Sample Output: show route receive-protocol bgp extensive on page 776

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Nexthop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

AS path: I <Originator>—(For route reflected output only) Route reflector set the originator ID attribute.

Cluster list—(For router reflected output only) Cluster id sent by the route reflector.

Originator ID—(For router reflected output only) Address of router that originally sent route to the route reflector.

Route Distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.

VPN Label—Label assigned to this VPN.

Communities—Community path attribute of the route.

Sample Output: show route receive-protocol bgp detail

```
 user@host> show route receive-protocol bgp 10.255.14.171 detail

inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY, control
flags:
0, mtu: 0

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY, control
flags:0, mtu: 0
```

Sample Output: show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive

inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path

frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY, control
flags:0, mtu: 0

bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY, control
flags:0, mtu: 0
```

show route table

Syntax show route table *routing-table*

Description Display the route entries in a particular routing table.

Options *routing-table*—Name of a routing table, in the format inet.*number*.

Required Privilege Level view

Sample Output Sample Output: show route table on page 778
Sample output: show route table detail | find on page 778
Sample output: show route table bgp.l2.vpn on page 778

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

label stacking—On the penultimate-hop router, indicates the depth of the MPLS label stack, where the label popping operation is needed. A pair of routes is displayed because the pop operation is different depending on the stack depth.

- S=0 route indicates the packet with an incoming label stack depth n >=2 exits this router using MPLS, with *n*–1 stack label.
- No S= information indicates a normal MPLS route, which has a stack depth of 1.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

Sample Output: show route table

```
user@host>show route table vpn-a

vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1:1:1:1/96
                   *[VPN/7] 05:48:27
                     Discard
192.168.24.1:1:2:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

Sample output: show route table detail | find

```
user@host> show route table vpn-a detail | find 192.168.24.1:1:2:1/96
192.168.24.1:1:2:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 192.168.24.1:1
                Source: 192.168.24.1
                Nexthop: 10.0.16.2 via fe-0/0/1.0, selected
                label-switched-path am
                Push 100067
                Protocol Nexthop: 192.168.24.1 Indirect nexthop: 84cfc38 39
                State: <Secondary Active Int Ext>
                Local AS: 65412 Peer AS: 65412
                Age: 4:56       Metric2: 3
                Task: BGP_65412.192.168.24.1+1028
                Announcement bits (1): 0-vpn-a-OSPF
                AS path: I
                Communities: target:65412:200 Layer2-info: encaps:VLAN, control
                flags:0, mtu: 0
```

Sample output: show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

Chapter 46

Layer 3 VPN Monitoring and Troubleshooting

Table 98 summarizes the command-line interface (CLI) commands you can use to monitor Layer 3 VPNs. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

These commands have been condensed and grouped to specifically address VPNs, with references to the complete command elsewhere in this book.

Table 98: Commands for Monitoring Layer 3 VPNs

Task Category	Task or Information to Monitor	Command
BGP Monitoring Information	Entries in the BGP neighbor database.	show bgp neighbor on page 781
	BGP summary information.	show bgp summary on page 787
Forwarding Table Information	Information about the entries in the kernel's forwarding table.	show route forwarding-table on page 797
OSPF Monitoring Information	Link-state database entries.	show ospf database on page 789
	Status of interfaces on which OSPF is running.	show ospf interface on page 790
	SPF log.	show ospf log on page 790
	OSPF neighbor.	show ospf neighbor on page 791
	OSPF routing table entries.	show ospf route on page 791
	General OSPF statistics.	show ospf statistics on page 792
RIP Monitoring Information	Display information about all RIP neighbors.	show rip neighbor on page 793
	Display rip protocol statistics.	show rip statistics on page 793

Task Category	Task or Information to Monitor	Command
Routing Table Information	Information about the entries in the routing tables.	show route on page 794
	Routes transmitted by a particular routing protocol.	show route advertising-protocol on page 795
	Route instance information.	show route instance on page 800
	Routes learned by the specified protocol.	show route protocol on page 802
	Routes received by a particular routing protocol.	show route receive-protocol on page 805
	Routes in a particular routing table.	show route table on page 809
System Management	Check the reachability of a remote VPN node.	ping vpn-interface on page 780
	Open secure shell connections between the local router and a remote system.	ssh on page 812
	Open a telnet session to a remote system.	telnet on page 813
	Determine the route to a host.	traceroute on page 814

ping vpn-interface

Syntax

ping vpn-interface *vpn-interface host* <local *echo-address*>

Description

Check the reachability of a remote vpn-interface router.

The ping vpn-interface command provides the source address of the ICMP packets and refers to the appropriate VPN routing and forwarding (VRF) table to perform a route lookup.

Type Ctrl-C to interrupt a ping vpn-interface command.

Options

host—IP address or name of remote system to ping.

local *echo-address*—(Optional) Set the source IP address (the local address) in outgoing ECHO_REQUEST packets.
If default-address-selection is configured, lo0 is the source.
If default-address-selection is not configured, outgoing interface is the source.

vpn-interface—Name of VPN interface from which to ping. Specify address including the logical interface. If no logical interface is specified, lo0 is used.

Sample Output

```
user@vpnhost> ping vpn-interface fe-1/1/0.1 192.168.197.94 local 192.168.197.93
count 3
PING 192.168.197.94 (192.168.197.94): 56 data bytes
64 bytes from 192.168.197.94: icmp_seq=0 ttl=255 time=0.866 ms
64 bytes from 192.168.197.94: icmp_seq=1 ttl=255 time=0.728 ms
64 bytes from 192.168.197.94: icmp_seq=2 ttl=255 time=0.753 ms
--- 192.168.197.94 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss round-trip
min/avg/max/stddev = 0.728/0.782/0.866/0.060 ms
```

show bgp neighbor

Syntax show bgp neighbor <*neighbor-address*>

Description Display information about BGP peers.

Options none—Display information about all BGP peers.

neighbor-address—(Optional) IP address of a BGP peer.

Required Privilege Level view

See Also For information about the local-address, nlri, hold-time, and preference statements, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols.*

Output Fields Peer—Address of the BGP neighbor. The address is followed by the neighbor's port number.

AS—AS number of the peer.

Local—Address of the local router. The address is followed by the peer's port number.

Type—Type of peer. It can be either Internal or External.

State—Current state of the BGP session. It can be one of the following:

- Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message.
- Connect—BGP is waiting for the transport protocol connection to complete.
- Established—The BGP session has been established, and the peers are exchanging update messages.
- Idle—This is the first stage of a connection. BGP is waiting for a Start event.
- OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message.
- OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.

Flags—Internal BGP flags. It can be one or more of the following:

- CleanUp—The peer session is being shut down.
- Delete—This peer has been deleted.
- Idled—This peer has been permanently idled.
- Initializing—The peer session is initializing.
- SendRtn—Messages are being sent to the peer.
- TryConnect—Another attempt is being made to connect to the peer.

- Unconfigured—This peer is not configured.
- WriteFailed—An attempt to write to this peer failed.

Last State—Previous state of the BGP session.

Last Event—Last activity that occurred in the BGP session. It can be one of the following:

- Closed—The BGP session closed.
- ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect.
- HoldTime—The session ended because the hold timer expired.
- KeepAlive—The local router sent a BGP keepalive message to the peer.
- Open—The local router sent a BGP open message to the peer.
- OpenFail—The local router did not receive an acknowledgment of a BGP open message from the peer.
- RecvKeepAlive—The local router received a BGP keepalive message from the peer.
- RecvNotify—The local router received a BGP notification message from the peer.
- RecvOpen—The local router received a BGP open message from the peer.
- RecvUpdate—The local router received a BGP update message from the peer.
- Start—The peering session started.
- Stop—The peering session stopped.
- TransportError—A TCP error occurred.

Last Error—Last error that occurred in the BGP session. It can be one of the following:

- Cease—An error occurred, such as a version mismatch, that caused the session to close.
- Finite State Machine Error—In setting up the session, BGP received a message that it did not understand.
- Hold Time Expired—The session's hold time expired.
- Message Header Error—The header of a BGP message was malformed.
- Open Message Error—A BGP open message contained an error.
- None—No errors occured in the BGP session.
- Update Message Error—A BGP update message contained an error.

Export—Name of the export policy that is configured on the peer.

Import—Name of the import policy that is configured on the peer.

Options—Configured BGP options. It can be one or more of the following:

- AddressFamily—Configured address family. It can be inet or inet-vpn.
- GracefulRestart—Graceful restart is configured.
- HoldTime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.
- Local Address—Address configured with the local-address statement.
- NLRI—Configured MBGP state for the BGP group. It can be either multicast or unicast, or both if you have configured nlri any.
- Peer AS—Configured peer AS.
- Preference—Preference value configured with the preference statement. The default preference value is 170.
- Refresh—Configured to refresh automatically when policy changes.
- Rib-group—Configured routing table group.

Address families configured—Names of configured address families for VPN.

Local Address—The interface used by an external BGP session (except when you use the BGP multihop option for this peer).

Holdtime—Hold time configured with the hold-time statement. The default hold time is 90 seconds. The hold time is three times the interval at which keepalive messages are sent.

Preference—Preference value configured with the preference statement. The default preference value is 170.

Number of flaps—Number of times the BGP session has gone down and then come back up.

Peer ID—Router identifier of the peer.

Local ID—Router identifier of the local router.

Active hold time—Hold time the local router negotiated with the peer.

Keepalive interval—Keepalive interval, in seconds.

Local Interface—Name of directly connected interface.

NLRI for restart configured on peer—Names of address families configured for restart.

NLRI advertised by peer—Address families supported by the peer. It can be unicast or multicast.

NLRI for this session—Address families being used for this session.

Peer supports Refresh capability—Remote peer supports Route Refresh Capability, and is able to send and request full route table readvertisements. For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.

Restart time configured on the peer—Configured time allowed for restart on the neighbor. The range is 1-600 seconds. The default value is 300 seconds.

Stale routes from peer are kept for—When graceful restart is negotiated, this is the maximum time allowed to hold routes from neighbors after the BGP session has gone down. The range is 1-600 seconds. The default value is 300 seconds.

Restart time requested by this peer—Restart time requested by this neighbor during capability negotiation.

NLRI that peer supports restart for—Address families indicates graceful restart during capability negotiation.

NLRI that restart is negotiated for—Both address families are supported by both this peer and its neighbor.

NLRI of received end-of-rib markers—Address families for which end-of-rib markers are received from the neighbor.

NLRI of all end-of-rib markers sent—Address families for which end-of-rib markers are sent to the neighbor.

Table inet.*number*—Information about the routing table.

- RIB State—BGP is in the graceful restart process for this routing table. It can be restart is complete or restart in progress.
- Bit—Number that represents the entry in the routing table for this peer.
- Send state—State of the BGP group. It can be in sync, not in sync, or not advertising.
- Active prefixes—Number of prefixes received from the peer that are active in the routing table.
- Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table.
- Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

Last traffic (seconds)—Last time any traffic was received from the peer or sent to the peer, and the last time the local router checked.

Input messages—Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output messages—Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, and the buffer size in octets. The buffer size is 16 KB.

Output queue[0], Output queue[1]—Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.

Trace options—Configured tracing of BGP protocol packets and operations. See the *JUNOS Internet Software Configuration Guide: Getting Started*.

Trace file—Name of the file to receive the output of the tracing operation.

Sample Output

```
user@host> show bgp neighbor
Peer: 10.39.1.5+2135  AS 2     Local: 10.39.1.6+179   AS 1
        Type: External    State: Established    Flags: <>
        Last State: OpenConfirm   Last Event: RecvKeepAlive
        Last Error: None
        Export: [ ospf-to-bgp stat-to-bgp ]
        Options: <Preference LocalAddress HoldTime AddressFamily PeerAS
      Rib-group Refresh>
        Address families configured: inet-unicast
        Local Address: 10.39.1.6 Holdtime: 90 Preference: 170
        Number of flaps: 0
        Peer ID: 10.255.245.245   Local ID: 10.255.71.14     Active Holdtime: 90

        Keepalive Interval: 30
        NLRI advertised by peer: inet-unicast
        NLRI for this session: inet-unicast
        Peer supports Refresh capability (2)
        Table VPN-AB.inet.0 Bit: 30000
          Send state: in sync
          Active prefixes: 1
          Received prefixes: 1
          Suppressed due to damping: 0
        Last traffic (seconds): Received 26   Sent 25   Checked 25
        Input messages:  Total 22     Updates 8       Refreshes 0     Octets 714

        Output messages: Total 23     Updates 8       Refreshes 0     Octets 737

        Output Queue[2]: 0


      Peer: 10.255.71.15+1619 AS 1     Local: 10.255.71.14+179 AS 1
        Type: Internal    State: Established    Flags: <>
        Last State: OpenConfirm   Last Event: RecvKeepAlive
        Last Error: None
        Options: <Preference LocalAddress HoldTime AddressFamily Rib-group
      Refresh>
        Address families configured: inet-vpn-unicast
        Local Address: 10.255.71.14 Holdtime: 90 Preference: 170
        Number of flaps: 0
        Peer ID: 10.255.71.15     Local ID: 10.255.71.14     Active Holdtime: 90

        Keepalive Interval: 30
        NLRI advertised by peer: inet-vpn-unicast
        NLRI for this session: inet-vpn-unicast
        Peer supports Refresh capability (2)
        Table bgp.l3vpn.0 Bit: 10001
          Send state: in sync
          Active prefixes: 2
```

```
  Received prefixes: 2
  Suppressed due to damping: 0
Table VPN-A.inet.0 Bit: 20001
  Send state: in sync
  Active prefixes: 1
  Received prefixes: 1
  Suppressed due to damping: 0
Table VPN-AB.inet.0 Bit: 30001
  Send state: in sync
  Active prefixes: 2
  Received prefixes: 2
  Suppressed due to damping: 0
Table VPN-B.inet.0 Bit: 40001
  Send state: in sync
  Active prefixes: 1
  Received prefixes: 1
  Suppressed due to damping: 0
Last traffic (seconds): Received 26   Sent 25   Checked 25
Input messages:  Total 20     Updates 6       Refreshes 0     Octets 934

Output messages: Total 22     Updates 5       Refreshes 2     Octets 798

Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
```

show bgp summary

Syntax show bgp summary

Description Display summary information about BGP and its neighbors.

Required Privilege Level view

Output Fields Groups—Number of BGP groups.

Peers—Number of BGP peers.

Down peers—Number of down BGP peers.

Table—Name of routing table.

Tot Paths—Total number of paths.

Act Paths—Number of active routes.

Suppressed—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

History—Number of withdrawn routes stored locally to keep track of damping history.

Damp State—Number of routes with a Figure of Merit (FOM) greater than 0, but still active because the FOM has not reached the threshold at which suppression occurs.

Pending—Routes in process by BGP import policy.

Peer—Address of each BGP peer. Each peer has one line of output.

AS—Peer's AS number.

InPkt—Number of packets received from the peer.

OutPkt—Number of packets sent to the peer.

OutQ—Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.

Flaps—Number of times the BGP session has gone down and then come back up.

Last Up/Down—Last up/down time since the neighbor transitioned to or from the established state.

State/#Active/Received/Damped—Displays either the BGP state or, if the neighbor is in the established state, the number of paths received from the neighbor, as well as the number of these paths that have been accepted as active and hence are being used for forwarding.

Established—The BGP session has been established, and the peers are exchanging update messages.

Sample Output

```
user@host> show bgp summary
      Groups: 2 Peers: 2 Down peers: 0
      Table          Tot Paths  Act Paths Suppressed    History Damp State
      Pending
      bgp.l3vpn.0            2          2          0          0          0
         0
      Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
      State|#Active/Received/Damped...
      10.39.1.5           2         21         22       0       0        6:26
      Establ
        VPN-AB.inet.0: 1/1/0
      10.255.71.15        1         19         21       0       0        6:17
      Establ
        bgp.l3vpn.0: 2/2/0
        VPN-A.inet.0: 1/1/0
        VPN-AB.inet.0: 2/2/0
        VPN-B.inet.0: 1/1/0 d
```

show ospf database

Syntax show ospf database <summary | brief | detail | extensive> <instance *instance name*><*lsa-filters*>

Description Display the entries in the OSPF link-state database, which contains data about LSA packets.

Options brief—(Optional) Display a brief listing of all entries in the OSPF link-state database.

detail—(Optional) Display detailed information about the entries in the OSPF link-state database.

extensive—(Optional) Display extremely detailed information about the entries in the OSPF link-state database.

instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

lsa-filters—(Optional) One or more of the following LSA filters. If you specify more than one filter, only LSAs that match all the filters are displayed. For example, the command show ospf database detail router lsa-id 10.0.0.1 displays all router LSAs in all areas that have an LSA identifier of 10.0.0.1.

- advertising-router *address*—(Optional) Display the LSAs advertised by a particular router.
- area *area-id*—(Optional) Display the LSAs in a particular area.
- lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.
- *lsa-type*—(Optional) Display specific types of LSAs. You can specify asbrsummary, extern, netsummary, network, nssa, or router.

summary—(Optional) Display summary information about OSPF link-state database.

Default: brief

Required Privilege Level view

Output Fields See the output fields for show ospf database on page 596.

Sample Output See the sample output for show ospf database on page 596.

show ospf interface

Syntax show ospf interface <instance *instance name*> <brief | detail | extensive>

Description Display the status of interfaces that are running OSPF between the PE and CE routers.

Options brief—(Optional) Display brief status information about OSPF interfaces.

detail—(Optional) Display detailed status information about OSPF interfaces.

extensive—(Optional) Display extremely detailed status information about OSPF interfaces, including information about what is queued on the interfaces and neighbors. When OSPF is idle, the output of the extensive option is identical to that of the detail option.

instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

Default: brief

Required Privilege Level view

Output Fields See the output fields for show ospf interface on page 602.

Sample Output See the sample output for show ospf interface on page 602

show ospf log

Syntax show ospf log <instance *instance name*>

Description Display the entries in the OSPF log of shortest-path-first (SPF) calculations.

Options instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

Required Privilege Level view

Output Fields See the output fields for show ospf log on page 606.

Sample Output See the sample output for show ospf log on page 606.

show ospf neighbor

Syntax show ospf neighbor <*neighbor-address*> <instance *instance name*><brief | detail | extensive>

Description Display information about all OSPF neighbors.

Options brief—(Optional) Display brief information about OSPF neighbors.

detail—(Optional) Display detailed information about OSPF neighbors.

extensive—(Optional) Display extremely detailed information about OSPF neighbors.

instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

neighbor-address—(Optional) Address or ID of specific OSPF neighbor.

Default: brief

Required Privilege Level view

Output Fields See the output fields for show ospf neighbor on page 607.

Sample Output See the sample output for show ospf neighbor on page 607.

show ospf route

Syntax show ospf route <detail> <(abr | asbr | extern | inter | intra)><instance *instance name*>

Description Display the entries in the OSPF routing table.

Options none—Display all entries in the OSPF routing table.

abr—(Optional) Display OSPF routes to area border routers.

asbr—(Optional) Display OSPF routes to autonomous system border routers.

detail—(Optional) Display detailed information about the entries in the OSPF routing table.

extern—(Optional) Display external OSPF routes.

inter—(Optional) Display interarea OSPF routes.

intra—(Optional) Display intra-area OSPF routes.

instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

Required Privilege Level view

Output Fields See the output fields for show ospf route on page 610.

Sample Output See the sample output for show ospf route on page 610.

show ospf statistics

Syntax show ospf statistics <instance *instance name*>

Description Display OSPF protocol statistics.

Options instance *instance name*—(Optional) Display all OSPF interfaces under the named routing instance.

Required Privilege Level view

Output Fields Packet type—Type of OSPF packet.

Total Sent/Total Received—Total number of packets sent and received.

Last 5 seconds Sent/Last 5 seconds Received—Total number of packets sent and received in the last 5 seconds.

LSAs retransmitted—Total number of link-state advertisements transmitted, and number retransmitted in the last 5 seconds.

Receive errors—Number and type of receive errors.

Output Fields See the output fields for show ospf statistics on page 612.

Sample Output See the sample output for show ospf statistics on page 612.

show rip neighbor

Syntax show rip neighbor <*name*>

Description Display information about all RIP neighbors.

Options none—Display brief information about all RIP neighbors.

name—(Optional) Display detailed information about a specific RIP neighbor.

Required Privilege Level view

Output Fields Neighbor—Name of RIP neighbor.

State—State of the connection. The interface can be either Up or Dn (Down).

Source Address—Source address.

Destination Address—Destination address.

Send Mode—Send options. It could be broadcast, multicast, none, or version 1.

Receive Mode—Type of packets to accept. It could be both, none, version 1, or version 2.

In Met—Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output: show rip neighbor

```
user@host> show rip neighbor
                          Source          Destination     Send   Receive   In
Neighbor          State   Address         Address         Mode   Mode     Met
--------          -----   -------         -----------     ----   -------  ---
so-0/2/1.0           Up   10.39.1.25      224.0.0.9       mcast  both       1
```

show rip statistics

Syntax show rip statistics <*name*>

Description Display RIP protocol statistics.

Options none—Display brief information about all RIP neighbors.

name—(Optional) Display detailed information about a specific RIP neighbor.

Required Privilege Level view

Output Fields See the output fields for show rip statistics on page 627.

Sample Output See the sample output for show rip statistics on page 627.

show route

Syntax show route <*destination-prefix*> <all | hidden> <brief | detail | extensive> <*filters*>

Description Display the active entries in the routing tables.

Options none—Display all active entries in the routing tables.

all—(Optional) Display information about all routes, including hidden entries.

brief—(Optional) Display brief route information.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

detail—(Optional) Display detailed route information.

extensive—(Optional) Display very detailed route information.

filters—(Optional) One or more filters that select specific information to display from the routing tables. To simplify the readability and explanation of the filters, each filter is treated as an individual command even though you can combine filters in a single show route command. The following is the syntax of the show route command including all filters:

```
show route <destination-prefix> <summary | detail | extensive> <best | exact | range>
   <advertising-protocol protocol neighbor-address> <aspath-regex regular-expression>
   <community as-number:community-value> <damped> <detail> <forwarding-table>
   <inactive> <next-hop address> <output interface-name> <protocol protocol> <range>
   <receive-protocol protocol neighbor-address> <source-gateway address>
   <table inet.number>
```

hidden—(Optional) Display hidden route information.

Default: brief

Required Privilege Level view

Output Fields See the output fields for show route on page 516.

Sample Output See the sample output for show route on page 516.

show route advertising-protocol

Syntax show route advertising-protocol *protocol neighbor-address* <detail>

Description Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol. The information displayed reflects the routes that the routing table has exported into the routing protocol and that were filtered by that protocol's export routing policy statements. (In the figure "Importing and Exporting Routing Policies" in the section "How to Construct Routing Policy" in the *JUNOS Internet Software Configuration Guide: Policy Framework*, these are the routes shown in the right "Protocol" box.)

Options detail—(Optional) Display additional information that was advertised to the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were advertised, including communities, route reflection, and aggregator.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route. It can be bgp, dvmrp, msdp, pim, or rip.

Required Privilege Level view

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Nexthop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

Route-distinguisher—The unique route distinguisher associated with this routing instance.

VPN Label—Name of label assigned to this VPN.

Communities—Community path attribute of the route.

Sample Output: show route advertising-protocol (standard)

```
user@host> show route advertising-protocol bgp 10.255.14.171
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.172/32   Self                     1        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.181/32   Self                     2        100 I
```

Sample Output: show route advertising-protocol detail

```
user@host> show route advertising-protocol bgp 10.255.14.171 detail
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.172/32 (1 entry, 1 announced)
       BGP group ibgp type Internal
           Route Distinguisher: 10.255.14.174:100
           VPN Label: 100004
           Nexthop: Self
           MED: 1
           Localpref: 100
           AS path: I
           Communities: target:1:100 Route-Type:0.0.0.0:1:0
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.181/32 (1 entry, 1 announced)
       BGP group ibgp type Internal
           Route Distinguisher: 10.255.14.174:200
           VPN Label: 100005
           Nexthop: Self
           MED: 2
           Localpref: 100
           AS path: I
           Communities: target:1:200 Route-Type:0.0.0.0:1:0
```

show route forwarding-table

Syntax show route forwarding-table <*destination-prefix*> <family *family* | matching *matching*> <multicast><vpn *vpn*><detail | standard | summary>

Description Display the route entries in the kernel's forwarding table. This is the version of the forwarding table in the Routing Engine. The Routing Engine copies this table to the Packet Forwarding Engine.

The show route forwarding-table command displays only the network-layer prefixes and their next hops. Commonly, you need information in the show route commands. However, when you run into problems, such as connectivity problems, it can be necessary to use the show route forwarding-table command to verify that the routing protocol process has relayed the correct information into the forwarding table.

> **Note** The show route forwarding-table command is an independent command, not a filter that selects specific information that is displayed from the routing tables. You cannot use this command in conjunction with any of the show route filter options (see Table 98, "Commands for Monitoring Layer 3 VPNs" on page 779).

Options none—Display the routes in the forwarding table. This option is equivalent to the UNIX netstat –rn command.

destination-prefix—(Optional) IP address of a destination.

detail—(Optional) Display detailed information, including clone routes. This option is equivalent to the UNIX netstat –rna command.

family *family*—(Optional) Display routing table entries for the specified family.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

standard—(Optional) Display standard information about routes.

summary—(Optional) Display summary information about routes.

vpn *vpn*—(Optional) Display routing table entries for specified VPN.

Required Privilege Level view

Output Fields Internet, ISO, MPLS—IP, ISO, and MPLS addresses.

Destination—Destination of the route.

Type—How the route was placed into the forwarding table (located in the second column):

- clon—Clone route (for TCP or multicast only; displayed only if you specify the detail option).
- dest—Remote addresses directly reachable through an interface.

- iddn—dest route for which the interface is down.
- ifcl—clon route for which the interface is down.
- ifdn—intf route for which the interface is down.
- ignr—Ignore this route.
- intf—Installed as the result of configuring an interface.
- perm—Permanent (installed by kernel when the routing table is initialized).
- user—Installed by the routing protocol process or as a result of the configuration.

RtRef—IP, ISO, and MPLS addresses.

Inlf—Incoming interface.

Flags—Additional information about a route.

- 0x01—Static route.
- 0x02—Cache route.
- 0x04—Check against incoming interface.
- 0x08—Route has accounting.
- 0x10—Route has been sent to the Packet Forwarding Engine.

Nexthop—Next hop to the destination.

Type—Information about the next hop (located in the fourth column):

- bcst—Broadcast.
- dscd—Discard; no ICMP unreachable message sent.
- hold—Next hop is waiting to be resolved (will turn into a unicast or multicast).
- locl—Local address on an interface.
- mcrt—“Regular” multicast next hop.
- mcst—Wire multicast next hop (limited to the LAN).
- mdsc—Multicast discard.
- mgrp—Multicast group member.
- recv—Receive.
- rjct—Discard; ICMP unreachable message sent.
- rslv—Resolving next hop.

- ucst—Unicast.
- ulst—List of unicast next hops. A packet sent to this next hop will go to any next hop in the list.

Index—Number associated with the next hop.

NhRef—Number of routes that reference this next hop.

Netif—Interface used to reach the next hop.

indr—An index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.

Sample Output: show route forwarding-table vpn

```
user@host> show route forwarding-table vpn VPN-A
      Routing table:: VPN-A.inet
      Internet:
      Destination        Type RtRef Nexthop           Type Index NhRef Netif
      default            perm     0                   dscd     4     4
      10.39.10.20/30     intf     0 ff.3.0.21         ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21       locl    36     1
      10.255.14.172/32   user     0                   ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                   indr    81     3
                                                      Push 100004, Push
      100004(top) so-1/0/0.0
      224.0.0.0/4        perm     2                   mdsc     5     3
      224.0.0.1/32       perm     0 224.0.0.1         mcst     1     8
      224.0.0.5/32       user     1 224.0.0.5         mcst     1     8
      255.255.255.255/32 perm     0                   bcst     2     3
```

show route instance

Syntax show route instance <detail | summary><*routing-instance-name*>

Description Display routing instance information.

Options detail—(Optional) Display detailed routing instance information.

routing-instance-name—(Optional) Display routing instance information for specified instance.

summary—(Optional) Display summary routing instance information.

Required Privilege Level view

Output Fields Instance—Name of routing instance.

Type—Type of routing instance. It can be forwarding, no-forwarding, or vrf.

State—State of the routing instance. It can be active or inactive.

Interfaces—Name of interfaces belonging to this routing instance.

Route-distinguisher—The unique route distinguisher associated with this routing instance.

Vrf-import—VPN routing and forwarding instance import policy name.

Vrf-export—VPN routing and forwarding instance export policy name.

Tables—Tables associated with this routing instance.

Primary rib—Primary table for this routing instance.

Active/holddown/hidden—Number of active, holddown, and hidden routes.

Sample Output: show route instance detail

```
user@host> show route instance detail
yellow-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/3.0
  Route-distinguisher: 69:40
  Vrf-import: [ yellow-vpn-import ]
  Vrf-export: [ yellow-vpn-export ]
  Tables:
    yellow-vp.inet.0       : 5 routes (5 active, 0 holddown, 0 hidden)
pink-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/1.0
  Route-distinguisher: 69:30
  Vrf-import: [ pink-vpn-import ]
  Vrf-export: [ pink-vpn-export ]
  Tables:
    pink-vpn.inet.0        : 4 routes (4 active, 0 holddown, 0 hidden)
```

```
green-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/2.0
  Route-distinguisher: 69:20
  Vrf-import: [ green-vpn-import ]
  Vrf-export: [ green-vpn-export ]
  Tables:
    green-vpn.inet.0       : 6 routes (6 active, 0 holddown, 0 hidden)
blue-vpn:
  Type: vrf               State: Active
  Interfaces:
    fe-1/0/0.0
  Route-distinguisher: 69:10
  Vrf-import: [ blue-vpn-import ]
  Vrf-export: [ blue-vpn-export ]
  Tables:
    blue-vpn.inet.0        : 5 routes (5 active, 0 holddown, 0 hidden)
master:
  Type: forwarding        State: Active
  Tables:
    inet.0                 : 30 routes (29 active, 0 holddown, 1 hidden)
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    mpls.0                 : 9 routes (9 active, 0 holddown, 0 hidden)
```

Sample Output: show route instance summary

```
user@host> show route instance summary
Instance           Type                 Primary rib      Active/holddown/hidden
yellow-vpn         vrf                  yellow-vp.inet.0 5/0/0
pink-vpn           vrf                  pink-vpn.inet.0  4/0/0
green-vpn          vrf                  green-vpn.inet.0 6/0/0
blue-vpn           vrf                  blue-vpn.inet.0  5/0/0
master             forwarding           inet.0           29/0/1
```

show route protocol

Syntax show route protocol *protocol*

Description Display the route entries in the routing table that were learned from a particular protocol.

Options *protocol*—Protocol from which the route was learned. It can be one of the following: aggregate, bgp, direct, local, mpls, ospf, rip, or static.

Required Privilege Level view

Sample Output
Sample Output: show route protocol bgp on page 803
Sample Output: show route protocol ospf on page 803
Sample Output: show route protocol ospf detail on page 804
Sample Output: show route protocol rip on page 804
Sample Output: show route protocol rip detail on page 804

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

via—Interface to reach the next hop.

Sample Output: show route protocol bgp

```
user@host> show route protocol bgp
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.173/32   *[BGP/170] 00:11:02, MED 1, localpref 100, from 10.255.14.172
                      AS path: I
                    >   via t3-0/2/1.0, Push 100163, Push 100009(top)

bgp.l3vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.180:200:10.255.14.173/32
                   *[BGP/170] 20:31:25, MED 1, localpref 100, from 10.255.14.172
                      AS path: I
                    >   via t3-0/2/1.0, Push 100163, Push 100009(top)
```

Sample Output: show route protocol ospf

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
224.0.0.5/32       *[OSPF/10] 20:25:55, metric 1


VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
224.0.0.5/32       *[OSPF/10] 20:26:20, metric 1
```

Sample Output: show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

10.255.14.173/32 (1 entry, 1 announced)
        *OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Active Int>
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                Announcement bits (2): 2-BGP.0.0.0.0+179 3-KRT
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0
```

Sample Output: show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both


VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32       *[RIP/100] 00:03:59, metric 1
```

Sample Output: show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both


VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route receive-protocol

Syntax show route receive-protocol *protocol neighbor-address* <detail | extensive | standard>

Description Display the routing information as it was received through a particular neighbor of a particular dynamic routing protocol. This information includes the routes that the local router advertised to the neighbor. The information displayed reflects the routes before they are filtered by that protocol's import routing policy statements and is placed into the routing table.

Options none—Display information that was received from the neighbor.

detail—(Optional) Display additional information that was received from the neighbor. For example, for BGP, this option displays information about any other BGP path attributes that were received, including communities, route reflection, and aggregator.

extensive—(Optional) Display extensive information that was received from the neighbor.

neighbor-address—Address of the neighboring router from which the route entry was received.

protocol—Protocol transmitting the route. It can be either bgp or rip.

Required Privilege Level view

Sample Output Sample Output: show route receive-protocol bgp (standard) on page 806
Sample Output: show route receive-protocol bgp detail on page 806
Sample Output: show route receive-protocol bgp extensive on page 808

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

Prefix—Route address.

Nexthop—Address of the next hop to the address.

MED—MED value included in the route.

Lclpref—Local preference value included in the route.

AS path—AS path included in the route.

Route Distinguisher—IP subnets are augmented with a 64-bit prefix called a route distinguisher to make them unique.

VPN Label—Label assigned to this VPN.

Communities—Community path attribute of the route.

AS path: I <Originator>—(For router reflected output only) Route reflector set the originator ID attribute.

Cluster list—(For router reflected output only) Cluster id sent by the route reflector.

Originator ID—(For router reflected output only) Address of router that originally sent route to the route reflector.

Sample Output: show route receive-protocol bgp (standard)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                      100 2 I
      10.255.14.179/32   10.255.14.171            2         100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                      100 2 I
      10.255.14.177/32   10.255.14.171                      100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                      100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171            2         100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                      100 2 I
```

Sample Output: show route receive-protocol bgp detail

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.175/32 (1 entry, 1 announced)
           Route Distinguisher: 10.255.14.171:200
           VPN Label: 100004
           Nexthop: 10.255.14.171
           Localpref: 100
           AS path: 2 I
           Communities: target:1:100 target:1:200
      10.255.14.179/32 (1 entry, 1 announced)
           Route Distinguisher: 10.255.14.171:100
           VPN Label: 100007
           Nexthop: 10.255.14.171
           MED: 2
           Localpref: 100
           AS path: I
           Communities: target:1:100 Route-Type:0.0.0.0:1:0
```

```
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
10.255.14.175/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.171:200
     VPN Label: 100004
     Nexthop: 10.255.14.171
     Localpref: 100
     AS path: 2 I
     Communities: target:1:100 target:1:200
10.255.14.177/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.171:300
     VPN Label: 100005
     Nexthop: 10.255.14.171
     Localpref: 100
 AS path: I
     Communities: target:1:200
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path

mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix             Nexthop                MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.171:300
     VPN Label: 100005
     Nexthop: 10.255.14.171
     Localpref: 100
     AS path: I
     Communities: target:1:200
10.255.14.171:100:10.255.14.179/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.171:100
     VPN Label: 100007
     Nexthop: 10.255.14.171
     MED: 2
     Localpref: 100
     AS path: I
     Communities: target:1:100 Route-Type:0.0.0.0:1:0
10.255.14.171:200:10.255.14.175/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.171:200
     VPN Label: 100004
     Nexthop: 10.255.14.171
     Localpref: 100
     AS path: 2 I
     Communities: target:1:100 target:1:200
```

Sample Output: show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop                MED    Lclpref AS path
     1.1.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     165.3.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     165.4.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     195.1.2.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68

     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop                MED    Lclpref AS path

     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route table

Syntax show route table *routing-table*

Description Display the route entries in a particular routing table.

Options *routing-table*—Name of a routing table, in the format inet.*number*.

Required Privilege Level view

Sample Output Sample Output: show route table (standard) on page 810
Sample Output: show route table detail on page 810
Sample Output: show route table on page 811

Output Fields *routing-table-name*—Name of the routing table; for example, inet.0.

destinations—Number of destinations for which there are routes in the routing table.

routes—Number of routes in the routing table:

- active—Number of routes that are active.
- holddown—Number of routes that are in the hold-down state prior to being declared inactive.
- hidden—Number of routes not used because of routing policy.

prefix—Route address.

label stacking—On the penultimate-hop router, indicates the depth of the MPLS label stack, where the label popping operation is needed. A pair of routes is displayed because the pop operation is different depending on the stack depth.

- S=0 route indicates the packet with an incoming label stack depth n >=2 exits this router using MPLS, with *n*–1 stack label.
- No S=information indicates a normal MPLS route, which has a stack depth of 1.

[*protocol*/*preference*]—Protocol from which the route was learned and the route's preference value.

+—A plus sign before [*protocol*/*preference*] indicates the active route, which is the route installed from the routing table into the forwarding table.

– —A hyphen before [*protocol*/*preference*] indicates the last active route.

*—An asterisk before [*protocol*/*preference*] indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.

time—How long the route has been known.

metric 1—First metric value in the route.

via—Interface to reach the next hop.

Sample Output: show route table (standard)

```
user@host> show route table bgp.l3vpn.0
       bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
       + = Active Route, - = Last Active, * = Both
       10.255.71.15:100:10.255.71.17/32
                           *[BGP/170] 00:03:59, MED 1, localpref 100, from
       10.255.71.15
                              AS path: I
                            > via so-2/1/0.0, Push 100020, Push 100011(top)
       10.255.71.15:200:10.255.71.18/32
                           *[BGP/170] 00:03:59, MED 1, localpref 100, from
       10.255.71.15
                              AS path: I
                            > via so-2/1/0.0, Push 100021, Push 100011(top)
```

Sample Output: show route table detail

```
user@host> show route table bgp.l3vpn.0 detail
       bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
       + = Active Route, - = Last Active, * = Both
       10.255.71.15:100:10.255.71.17/32 (1 entry, 0 announced)
               *BGP    Preference: 170/-101
                       Route Distinguisher: 10.255.71.15:100
                       Source: 10.255.71.15
                       Nexthop: via so-2/1/0.0, selected
                       Push 100020, Push 100011(top)
                       Protocol Nexthop: 10.255.71.15 Push 100020
                       Indirect nexthop: 84dcee0 125
                       State: <Active Int Ext>
                       Local AS:     1 Peer AS:     1
                       Age: 4:36       Metric: 1       Metric2: 1
                       Task: BGP_1.10.255.71.15+1619
                       AS path: I
                       Communities: target:1:100 Route-Type:0.0.0.0:1:0
                       VPN Label: 100020
                       Localpref: 100
                       Router ID: 10.255.71.15
                       Secondary tables: VPN-AB.inet.0 VPN-A.inet.0
       10.255.71.15:200:10.255.71.18/32 (1 entry, 0 announced)
               *BGP    Preference: 170/-101
                       Route Distinguisher: 10.255.71.15:200
                       Source: 10.255.71.15
                       Nexthop: via so-2/1/0.0, selected
                       Push 100021, Push 100011(top)
                       Protocol Nexthop: 10.255.71.15 Push 100021
                       Indirect nexthop: 84dce58 123
                       State: <Active Int Ext>
                       Local AS:     1 Peer AS:     1
                       Age: 4:53       Metric: 1       Metric2: 1
                       Task: BGP_1.10.255.71.15+1619
                       AS path: I
                       Communities: target:1:200 Route-Type:0.0.0.0:1:0
                       VPN Label: 100021
                       Localpref: 100
                       Router ID: 10.255.71.15
                       Secondary tables: VPN-B.inet.0 VPN-AB.inet.0
```

Sample Output: show route table

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both
      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                           Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                            AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```

ssh

Syntax ssh *host* < vpn-interface *vpn-interface*> <v2>

Description Use the secure shell (ssh) program to open a connection between a local router and a remote system and execute commands on the remote system. You can issue the ssh command from the JUNOS CLI to log into a remote system or from a remote system to log into the local router.

The ping vpn-interface command provides the source address to use for the sshsession and refers to the appropriate VPN routing and forwarding (VRF) table to perform a route lookup.

To open an SSH connection from a remote system to the router, use the ssh command on the remote system. When executing this command, you include one or more CLI commands by including them in quotation marks and separating the commands with semicolons:

ssh *address* '*cli-command1*; *cli-command2*'

To configure an ssh version 1 key for your user account, include the authentication ssh-rsa statement at the [edit system login user *user-name*] hierarchy level, as described in the *JUNOS Internet Software Configuration Guide: Getting Started*.

To confure an ssh version 2 key for your user account, include the authentication dsa-rsa statement at the [edit system login user *user-name*] hierarchy level, as described in the *JUNOS Internet Software Configuration Guide: Getting Started*.

Options *host*—Name or address of the remote system.

v2—(Optional) Use ssh version 2 only when connecting to a remote host.

vpn-interface *vpn-interface*—(Optional) VPN interface for ssh session. This option does not work when default-address-selection is configured.

Required Privilege Level network

telnet

Syntax telnet *host* <8bit><port *port*> <vpn-interface *vpn-interface*> <source *source*>

Description Open a telnet session to a remote system.

The ping vpn-interface command provides the source address to use for the telnet session and refers to the appropriate VPN routing and forwarding (VRF) table to perform a route lookup.

Type Ctrl-] to escape from the telnet session to the telnet command level, and then type quit to exit from telnet.

Options 8bit—(Optional) Use an 8-bit data path.

host—Name or address of the remote system.

port *port*—(Optional) Port number or service name on remote system.

source *source*—(Optional) Source address of telnet connection.

vpn-interface *vpn-interface*—(Optional) VPN interface for SSH session. This option does not work when default-address-selection is configured.

Required Privilege Level network

Sample Output

```
user@host> telnet 192.154.1.254
Trying 192.154.169.254...
Connected to berry.juniper.net.
Escape character is '^]'.
ttypa
login: <Control>]
telnet> quit
Connection closed.
```

traceroute

Syntax traceroute *host* <gateway *address*> <ttl *value*> <wait *seconds*> <noresolve> <source *address*> <tos *value*> <vpn-interface *vpn-interface*>

Description Determine the route to a host.

Options *host*—IP address or name of remote host.

gateway *address*—(Optional) Address of a router through which the route should transit.

noresolve—(Optional) Do not attempt to determine the host name that corresponds to the IP address.

source *address*—(Optional) IP address to use as the source address in outgoing packets.

tos *value*—(Optional) Value to include in the IP type of service field.
Range: 0 through 255

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request.
Range: 0 through 128

vpn-interface *vpn-interface*—(Optional) Name of VPN interface over which to send packets.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

Sample Output: traceroute between CE routers

```
user@host> traceroute vpn09
traceroute to vpn09.englab.skybank.net (10.255.14.179), 30 hops max, 40
      byte packets
       1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
       2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
           MPLS Label=100006 CoS=0 TTL=1 S=1
       3  vpn09.englab.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

Part 8
Security

■ IPSec Monitoring on page 817

Chapter 47
IPSec Monitoring

Table 99 summarizes the command-line interface (CLI) commands you can use to monitor and troubleshoot security. In the table, the commands are grouped by functionality. In the remainder of this chapter, they are explained alphabetically.

Table 99: Commands for Monitoring Security

Task Category	Task or Information to Monitor	Command
IPSec	Clear IKE security associations.	clear ike security-associations on page 817
	IKE security association information.	show ike security-associations on page 819
	Clear IPSec security associations.	clear ipsec security-associations on page 818
	IPSec security association information.	show ipsec security-associations on page 822

clear ike security-associations

Syntax clear ike security-associations <*destination ip-address*>

Description Clear information about the current IKE security association. This command is valid for dynamic type security associations only.

Options none—Clear all IKE security associations.

destination ip-address—(Optional) Clear a particular IKE security association to this destination address.

Required Privilege Level view

clear ipsec security-associations

Syntax clear ipsec security-associations <*sa-name*>

Description Clear information about the current IPSec security association. A new security association is then created. This command is valid for dynamic type security associations only.

Options none—Clear all IPSec security associations.

sa-name—(Optional) Clear a particular security association.

Required Privilege Level view

Sample Output: clear ipsec security-associations (detail)

```
user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

  Direction: inbound, SPI: 242379418, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 22979 seconds
  Hard lifetime: Expires in 28739 seconds

  Direction: outbound, SPI: 368592771, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 22979 seconds
  Hard lifetime: Expires in 28739 seconds

user@host> clear ipsec security-associations

user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

  Direction: inbound, SPI: 1031597683, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 23037 seconds
  Hard lifetime: Expires in 28797 seconds

  Direction: outbound, SPI: 1618419878, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Soft lifetime: Expires in 23037 seconds
  Hard lifetime: Expires in 28797 seconds
```

show ike security-associations

Syntax show ike security-associations <*peer-address*> <brief | detail>

Description Display information about the Internet Key Exchange (IKE) security associations (sa).

Options brief—(Optional) Display brief IKE security association information.

detail—(Optional) Display detailed IKE security association information.

peer-address—(Optional) The remote end of the IKE negotiation.

Default: brief

Required Privilege Level view

Output Fields IKE peer—The remote end of the IKE negotiation.

Role—Part played in the IKE session. If you are triggering the IKE negotiation, then you are the initiator. If you are accepting the first IKE exchange packets, then you are the responder.

Remote Address—(standard) Responder's address.

State—State of the IKE security association. If the state is mature, the IKE security association is established, If the state is not matured, the IKE security association is in the process of negotiation.

Initiator cookie—When triggering the IKE negotiation, a random number is sent to the remote node.

Responder cookie—Remote mode generates its own random number and sends it back to the initiator as a verification that the packets were received.

> **Note** Of the numerous security services available, protection against denial of service is one of the most difficult to address. A "cookie" or anti-clogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some denial of service attempts (such as simple flooding with bogus IP source addresses). Absolute protection against denial of service is impossible, but this anti-clogging token makes it easier.

Exchange type—Specifies the number of messages in a IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. JUNOS software supports two types of exchanges:

- Main—The exchange is done with 6 messages. Main encrypts the payload, protecting the identity of the neighbor.
- Aggressive—The exchange is done with 3 messages. Aggressive does not encrypt the payload, leaving the identity of the neighbor unprotected.

Authentication method—Type of authentication determines which payloads are exchanged and when they are exchanged. pre-shared keys is the only authentication method supported at this time.

Local—Prefix and port number of the local end.

Remote—Prefix and port number of the remote end.

Lifetime—Number of seconds remaining until the IKE security association expires.

Algorithms—Header for the IKE algorithms output.

Authentication—(Detail output only) Type of authentication algorithm used. It can be md5 or sha1.

Encryption—(Detail output only) Type of encryption algorithm used. It can be des-cbc, 3des-cbc, or None.

Pseudo random function— Function that generates highly unpredictable random numbers. It can be hmac-md5 or hmac-sha1.

Traffic statistics—Number of bytes and packets received and transmitted on the IKE security association.

Input bytes, Output bytes—Number of bytes received and transmitted on the IKE security association.

Input packets, Output packets—Number of packets received and transmitted on the IKE security association.

Flags—Notification to the key management daemon of the status of the IKE negotiation. It can be one of the following:

- caller notification sent—Caller program notified about the completion of the IKE negotiation.
- waiting for done—Negotiation is done. The library is waiting the for the remote end retransmission timers to expire.
- waiting for remove—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
- waiting for policy manager—Negotiation is waiting for a response from the policy manager.

IPSec security associations—Number of IPSec security associations created and deleted with this IKE security association.

Phase 2 negotiations in progress—Number phase 2 IKE negotiations in progress.

Negotiation type—Type of the phase 2 negotiation. At this time, JUNOS software supports quick mode.

Message ID—Unique identifier for a phase 2 negotiation.

Local identity—Identity of the local phase 2 negotiation. The format is
id-type-name(proto-name:port-number,[0..id-data-len]=iddata-presentation)

Remote identity—Identity of the remote phase 2 negotiation. The format is
id-type-name(proto-name:port-number,[0..id-data-len]=iddata-presentation)

Sample Output: show ike security-associations (standard)

```
user@host> show ike security-associations
Remote Address  State         Initiator cookie  Responder cookie
Exchange type
4.4.4.4         Matured       93870456fa000011  723a20713700003e  Main
```

Sample Output: show ike security-associations detail

```
user@host>show ike security-associations detail
IKE peer 4.4.4.4
  Role: Initiator, State: Matured
  Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 4.4.4.5:500, Remote: 4.4.4.4:500
  Lifetime: Expires in 187 seconds
  Algorithms:
   Authentication        : md5
   Encryption            : 3des-cbc
   Pseudo random function: hmac-md5
  Traffic statistics:
   Input  bytes  :                 1000
   Output bytes  :                 1280
   Input  packets:                    5
   Output packets:                    9
  Flags: Caller notification sent
  IPsec security assocations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
    Local: 4.4.4.5:500, Remote: 4.4.4.4:500
    Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
    Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
    Flags: Caller notification sent, Waiting for done
```

show ipsec security-associations

Syntax show ipsec security-associations <*sa-name*> <brief | detail>

Description Display information about the IPSec security associations applied to the local or transit traffic stream.

Options brief—(Optional) Display brief IPSec security association information.

detail—(Optional) Display detailed IPSec security association information.

sa-name—(Optional) Display a particular security association.

Default: brief

Required Privilege Level view

Sample Output Sample Output: show ipsec security-associations brief (with manual SA) on page 824
Sample Output: show ipsec security-associations detail (with manual SA) on page 824
Sample Output: show ipsec security-associations brief (with dynamic SA) on page 824
Sample Output: show ipsec security-associations detail (with dynamic SA) on page 824

Output Fields Security association—Name and interface family of the security association. If the interface family field is absent, it is a transport mode security association. The interface family can have one of three options:

- Up—The security association is referenced in the interface family and the interface family is up.
- Down—The security association is referenced in the interface family and the interface family is down.
- No reference—The security association is not referenced in the interface family.

Direction—Direction of the security association; it can be inbound or outbound.

SPI—Value of the security parameter index.

State—(Detail output only) State has two options, installed and not installed.

- Installed—The security association is installed in the security association database.
- Not installed—The security association is not installed in the security association database.

> **Note** For transport mode security associations, state should always be installed.

Mode—Mode of the security association. Mode can be transport or tunnel.

- transport—Protects single host-to-host protections.
- tunnel—Protects connections between security gateways.

Type—Type of the security association. Type can be manual or dynamic.

- manual—Security parameters require no negotiation. They are static, and are configured by the user.
- dynamic—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode.

Protocol—Protocol supported. transport mode supports Encapsulation Security Protocol (ESP) or Authentication Header (AH). tunnel mode supports ESP.

Authentication—(Detail output only) Type of authentication used. It can be hmac-md5-96, hmac-sha1-96, or none.

Encryption—(Detail output only) Type of encryption used. It can be des-cbc, 3des-cbc, or None.

Soft lifetime, Hard lifetime—(Detail and dynamic output only) Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The hard lifetime specifies the lifetime of the SA. The soft lifetime, which is derived from the hard lifetime, informs the IPSec key management system that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires.

- Expires in *seconds* seconds—The number of seconds left until the security association expires.
- Expires in *kilobytes* kilobytes—The number of kilobytes left until the security association expires.

Anti-replay service—State of the service that prevents packets from being replayed. It can be Enabled or Disabled.

Replay window size—The configured size of the anti-replay service window. The anti-replay window size protects the receiver against replay attacks by rejecting old or duplicate packets. It can be 32 or 64 packets. If the replay window size is 0, then the anti-replay service is disabled.

Sample Output: show ipsec security-associations brief (with manual SA)

```
user@host> show ipsec security-associations sa-manual brief
Security association: sa-manual, Interface family: No reference
  Direction SPI         Mode       Type     Protocol
  inbound   2312        tunnel     manual   ESP
  outbound  2312        tunnel     manual   ESP
```

Sample Output: show ipsec security-associations detail (with manual SA)

```
user@host> show ipsec security-associations sa-manual detail
Security association: sa-manual, Interface family: No reference

  Direction: inbound, SPI: 300, State: Not installed
  Mode: tunnel, Type: manual
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Anti-replay service: Disabled

  Direction: outbound, SPI: 300, State: Not installed
  Mode: tunnel, Type: manual
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
  Anti-replay service: Disabled
```

Sample Output: show ipsec security-associations brief (with dynamic SA)

```
user@host> show ipsec security-associations sa-dynamic brief
Security association: sa-dynamic, Interface family: Up
  Direction SPI         Mode       Type     Protocol
  inbound   3962975369  tunnel     dynamic  ESP
  outbound  3200053017  tunnel     dynamic  ESP
```

Sample Output: show ipsec security-associations detail (with dynamic SA)

```
user@host> show ipsec security-associations sa-dynamic detail
Security association: sa-dynamic, Interface family: Up

  Direction: inbound, SPI: 3962975369, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: des-cbc
  Soft lifetime: Expires in 20550 seconds
  Hard lifetime: Expires in 26310 seconds
  Anti-replay service: Enabled, Replay window size: 64

  Direction: outbound, SPI: 3200053017, State: Installed
  Mode: tunnel, Type: dynamic
  Protocol: ESP, Authentication: hmac-md5-96, Encryption: des-cbc
  Soft lifetime: Expires in 21990 seconds
  Hard lifetime: Expires in 26310 second
  Anti-replay service: Enabled, Replay window size: 64
```

Part 9
Appendix

■ Glossary on page 827

Appendix A
Glossary

A

AAL ATM adaptation layer. A series of protocols enabling various types of traffic, including voice, data, image, and video, to run over an ATM network.

active route Route chosen from all routes in the routing table to reach a destination. Active routes are installed into the forwarding table.

add/drop multiplexer *See ADM*.

Address Resolution Protocol *See ARP*.

adjacency Portion of the local routing information that pertains to the reachability of a single neighbor over a single circuit or interface.

ADM Add/drop multiplexer. SONET functionality that allows lower-level signals to be dropped from a high-speed optical connection.

aggregation Combination of groups of routes that have common addresses into a single entry in the routing table.

AH Authentication Header. A component of the IPSec protocol used to verify that the contents of a packet have not been changed, and to validate the identity of the sender. *See also ESP*.

ANSI American National Standards Institute. The United States' representative to the ISO.

APQ Alternate Priority Queuing. Dequeuing method that has a special queue, similar to SPQ, which is visited only 50 percent of the time. The packets in the special queue still have a predictable latency, although the upper limit of the delay is higher than that with SPQ. Since the other configured queues share the remaining 50 percent of the service time, queue starvation is usually avoided. *See also SPQ*.

APS Automatic Protection Switching. Technology used by SONET ADMs to protect against circuit faults between the ADM and a router and to protect against failing routers.

area Routing subdomain that maintains detailed routing information about its own internal composition and that maintains routing information that allows it to reach other routing subdomains. In IS-IS, an area corresponds to a Level 1 subdomain.

In IS-IS and OSPF, a set of contiguous networks and hosts within an autonomous system that have been administratively grouped together.

area border router Router that belongs to more than one area. Used in OSPF.

ARP Address Resolution Protocol. Protocol for mapping IP addresses to MAC addresses.

AS Autonomous system. Set of routers under a single technical administration. Each AS normally uses a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routers. Also called *routing domain*.

AS boundary router In OSPF, routers that exchange routing information with routers in other ASs.

AS external link advertisements OSPF link-state advertisement sent by AS boundary routers to describe external routes that they know. These link-state advertisements are flooded throughout the AS (except for stub areas).

AS path In BGP, the route to a destination. The path consists of the AS numbers of all routers a packet must go through to reach a destination.

ASIC Application-specific integrated circuit. Specialized processors that perform specific functions on the router.

ATM Asynchronous Transfer Mode. A high-speed multiplexing and switching method utilizing fixed-length cells of 53 octets to support multiple types of traffic.

atomic Smallest possible operation. An atomic operation is performed either entirely or not at all. For example, if machine failure prevents a transaction from completing, the system is rolled back to the start of the transaction, with no changes taking place.

Authentication Header *See AH*.

Automatic Protection Switching *See APS*.

autonomous system *See AS*.

autonomous system boundary router In OSPF, routers that exchange routing information with routers in other ASs.

autonomous system external link advertisements OSPF link-state advertisement sent by autonomous system boundary routers to describe external routes that they know. These link-state advertisements are flooded throughout the autonomous system (except for stub areas).

autonomous system path In BGP, the route to a destination. The path consists of the autonomous system numbers of all the routers a packet must pass through to reach a destination.

B

backbone area In OSPF, an area that consists of all networks in area ID 0.0.0.0, their attached routers, and all area border routers.

backplane On an M40 router, component of the Packet Forwarding Engine that distributes power, provides signal connectivity, manages shared memory on FPCs, and passes outgoing data cells to FPCs.

bandwidth The range of transmission frequencies a network can use, expressed as the difference between the highest and lowest frequencies of a transmission channel. In computer networks, greater bandwidth indicates faster data-transfer rate capacity.

Bellcore Bell Communications Research. Research and development organization created after the divestiture of the Bell System. It is supported by the regional Bell holding companies (RBHCs), which own the regional Bell operating companies (RBOCs).

BERT Bit error rate test. A test that can be run on a T3 interface to determine whether it is operating properly.

BGP Border Gateway Protocol. Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.

bit error rate test *See BERT*.

BITS Building Integrated Timing Source. Dedicated timing source that synchronizes all equipment in a particular building.

Border Gateway Protocol *See BGP*.

broadcast Operation of sending network traffic from one network node to all other network nodes.

bundle Collection of software that makes up a JUNOS software release.

C

CB Control Board. Part of the host subsystem that provides control and monitoring functions for router components.

CCC Circuit cross-connect. A JUNOS software feature that allows you to configure transparent connections between two circuits, where a circuit can be a Frame Relay DLCI, an ATM VC, a PPP interface, a Cisco HDLC interface, or an MPLS label-switched path (LSP).

CE device Customer edge device. Router or switch in the customer's network that is connected to a service provider's provider edge (PE) router and participates in a Layer 3 VPN.

CFM Cubic feet per minute. Measure of air flow in volume per minute.

Challenge Handshake Authentication Protocol *See CHAP*.

channel service unit *See CSU/DSU*.

CHAP A protocol that authenticates remote users. CHAP is a server-driven, three-step authentication mechanism that depends on a shared secret password that resides on both the server and the client.

CIDR Classless interdomain routing. A method of specifying Internet addresses in which you explicitly specify the bits of the address to represent the network address instead of determining this information from the first octet of the address.

CIP Connector Interface Panel. On an M160 router, the panel that contains connectors for the Routing Engines, BITS interfaces, and alarm relay contacts.

circuit cross-connect *See CCC*.

class of service *See CoS*.

CLEC (Pronounced “see-lek”) Competitive Local Exchange Carrier. Company that competes with the already established local telecommunications business by providing its own network and switching.

CLEI Common language equipment identifier. Inventory code used to identify and track telecommunications equipment.

CLI Command-line interface. Interface provided for configuring and monitoring the routing protocol software.

client peer In a BGP route reflection, a member of a cluster that is not the route reflector. *See also nonclient peer*.

CLNP Connectionless Network Protocol. ISO-developed protocol for OSI connectionless network service. CLNP is the OSI equivalent of IP.

cluster In BGP, a set of routers that have been grouped together. A cluster consists of one system that acts as a route reflector, along with any number of client peers. The client peers receive their route information only from the route reflector system. Routers in a cluster do not need to be fully meshed.

community In BGP, a group of destinations that share a common property. Community information is included as one of the path attributes in BGP update messages.

confederation In BGP, a group of systems that appears to external autonomous systems to be a single autonomous system.

constrained path In traffic engineering, a path determined using RSVP signaling and constrained using CSPF. The ERO carried in the packets contains the constrained path information.

core The central backbone of the network.

CoS Class of service. The method of classifying traffic on a packet-by-packet basis using information in the ToS byte to provide different service levels to different traffic.

CPE Customer premises equipment. Telephone or other service provider equipment located at a customer site.

craft interface Mechanisms used by a Communication Workers of America craftsperson to operate, administer, and maintain equipment or provision data communications. On a Juniper Networks router, the craft interface allows you to view status and troubleshooting information and perform system control functions.

CSCP Class Selector Codepoint.

CSNP Complete sequence number PDU. Packet that contains a complete list of all the LSPs in the IS-IS database.

CSPF Constrained Shortest Path First. An MPLS algorithm that has been modified to take into account specific restrictions when calculating the shortest path across the network.

CSU/DSU Channel service unit/data service unit. Channel service unit connects a digital phone line to a multiplexer or other digital signal device. Data service unit connects a DTE to a digital phone line.

customer edge device *See CE device*.

D

daemon Background process that performs operations on behalf of the system software and hardware. Daemons normally start when the system software is booted, and they run as long as the software is running. In the JUNOS software, daemons are also referred to as processes.

damping Method of reducing the number of update messages sent between BGP peers, thereby reducing the load on these peers without adversely affecting the route convergence time for stable routes.

data circuit-terminating equipment *See DCE*.

data-link connection identifier *See DLCI*.

data service unit *See CSU/DSU*.

Data Terminal Equipment *See DTE*.

dcd The JUNOS software interface process (daemon).

DCE Data circuit-terminating equipment. RS-232-C device, typically used for a modem or printer, or a network access and packet switching node.

default address Router address that is used as the source address on unnumbered interfaces.

denial of service *See DoS*.

dense wavelength-division multiplexing *See DWDM*.

designated router In OSPF, a router selected by other routers that is responsible for sending link-state advertisements that describe the network, which reduces the amount of network traffic and the size of the routers' topological databases.

destination prefix length Number of bits of the network address used for host portion of a CIDR IP address.

DHCP Dynamic Host Configuration Protocol. Allocates IP addresses dynamically so that they can be reused when they are no longer needed.

Diffie-Hellman A public key scheme, invented by Whitfield Diffie and Martin Hellman, used for sharing a secret key without communicating secret information, thus precluding the need for a secure channel. Once correspondents have computed the secret shared key, they can use it to encrypt communications.

Diffserv Differentiated Service (based on RFC 2474). Diffserv uses the ToS byte to identify different packet flows on a packet-by-packet basis. Diffserv adds a Class Selector Codepoint (CSCP) and a Differentiated Services Codepoint (DSCP).

Dijkstra algorithm *See SPF*.

DIMM Dual inline memory module. 168-pin memory module that supports 64-bit data transfer.

direct routes *See interface routes*.

DLCI Data-link connection identifier. Identifier for a Frame Relay virtual connection (also called a logical interface).

DoS Denial of service. System security breach in which network services become unavailable to users.

DRAM Dynamic random-access memory. Storage source on the router that can be accessed quickly by a process.

drop profile Drop probabilities for different levels of buffer fullness that are used by RED to determine from which queue to drop packets.

DSCP Differentiated Services Codepoint.

DSU Data service unit. A device used to connect a DTE to a digital phone line. Converts digital data from a router to voltages and encoding required by the phone line. *See also CSU/DSU*.

DTE Data Terminal Equipment. RS-232-C interface that a computer uses to exchange information with a serial device.

DVMRP Distance Vector Multicast Routing Protocol. Distributed multicast routing protocol that dynamically generates IP multicast delivery trees using a technique called reverse path multicasting (RPM) to forward multicast traffic to downstream interfaces.

DWDM Dense wavelength-division multiplexing. Technology that enables data from different sources to be carried together on an optical fiber, with each signal carried on its own separate wavelength.

Dynamic Host Configuration Protocol *See DHCP*.

E

EBGP External BGP. BGP configuration in which sessions are established between routers in different ASs.

ECSA Exchange Carriers Standards Association. A standards organization created after the divestiture of the Bell System to represent the interests of interexchange carriers.

edge router In MPLS, a router located at the beginning or end of a label-switching tunnel. When at the beginning of a tunnel, an edge router applies labels to new packets entering the tunnel. When at the end of a tunnel, the edge router removes labels from packets exiting the tunnel. *See also MPLS*.

EGP Exterior gateway protocol, such as BGP.

egress router In MPLS, last router in a label-switched path (LSP). *See also ingress router*.

EIA Electronic Industries Association. A United States trade group that represents manufacturers of electronics devices and sets standards and specifications.

EMI Electromagnetic interference. Any electromagnetic disturbance that interrupts, obstructs, or otherwise degrades or limits the effective performance of electronics or electrical equipment.

encapsulating security payload *See ESP*.

end system In IS-IS, network entity that sends and receives packets.

ERO Explicit Route Object. Extension to RSVP that allows an RSVP PATH message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing.

ESP Encapsulating security payload. A fundamental component of IPSec-compliant VPNs, ESP specifies an IP packet's encryption, data integrity checks, and sender authentication, which are added as a header to the IP packet. *See also AH*.

explicit path *See signaled path*.

Explicit Route Object *See ERO*.

export To place routes from the routing table into a routing protocol.

external BGP *See EBGP*.

external metric A cost included in a route when OSPF exports route information from external autonomous systems. There are two types of external metrics: Type 1 and Type 2. Type 1 external metrics are equivalent to the link-state metric; that is, the cost of the route, used in the internal autonomous system. Type 2 external metrics are greater than the cost of any path internal to the autonomous system.

F

fast reroute Mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP.

FEAC Far-end alarm and control. T3 signal used to send alarm or status information from the far-end terminal back to the near-end terminal and to initiate T3 loopbacks at the far-end terminal from the near-end terminal.

FEB Forwarding Engine Board. In M5 and M10 routers, provides route lookup, filtering, and switching to the destination port.

firewall A security gateway positioned between two different networks, usually between a trusted network and the Internet. A firewall ensures that all traffic that crosses it conforms to the organization's security policy. Firewalls track and control communications, deciding whether to pass, reject, discard, encrypt, or log them. Firewalls also can be used to secure sensitive portions of a local network.

FIFO First in, first out.

flap damping *See damping*.

flapping *See route flapping*.

Flexible PIC Concentrator *See FPC*.

Forwarding Engine Board *See FEB*.

forwarding information base *See forwarding table*.

forwarding table JUNOS software forwarding information base (FIB). The JUNOS routing protocol process installs active routes from its routing tables into the Routing Engine forwarding table. The kernel copies this forwarding table into the Packet Forwarding Engine, which is responsible for determining which interface transmits the packets.

FPC Flexible PIC Concentrator. An interface concentrator on which PICs are mounted. An FPC inserts into a slot in a Juniper Networks router. *See also PIC.*

FRU Field-replaceable unit. Router component that customers can replace onsite.

G

group A collection of related BGP peers.

H

hash A one-way function that takes a message of any length and produces a fixed-length digest. In security, a message digest is used to validate that the contents of a message have not been altered in transit. The Secure Hash Algorithm (SHA-1) and Message Digest 5 (MD5) are commonly used hashes.

Hashed Message Authentication Code *See HMAC*.

HDLC High-level data link control. An International Telecommunication Union (ITU) standard for a bit-oriented data link layer protocol on which most other bit-oriented protocols are based.

HMAC Hashed Message Authentication Code. A mechanism for message authentication that uses cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function—for example, MD5 or SHA-1—in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function.

hold time Maximum number of seconds allowed to elapse between the time a BGP system receives successive keepalive or update messages from a peer.

host module On an M160 router, provides routing and system management functions of the router. Consists of the Routing Engine and Miscellaneous Control Subsystem (MCS).

host subsystem Provides routing and system-management functions of the router. Consists of a Routing Engine and an adjacent Control Board (CB).

I

IANA Internet Assigned Numbers Authority. Regulatory group that maintains all assigned and registered Internet numbers, such as IP and multicast addresses. *See also NIC*.

IBGP Internal BGP. BGP configuration in which sessions are established between routers in the same ASs.

ICMP Internet Control Message Protocol. Used in router discovery, ICMP allows router advertisements that enable a host to discover addresses of operating routers on the subnet.

IDE Integrated Drive Electronics. Type of hard disk on the Routing Engine.

IEC International Electrotechnical Commission. *See ISO*.

IEEE Institute of Electronic and Electrical Engineers. International professional society for electrical engineers.

IETF Internet Engineering Task Force. International community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet.

IGMP Internet Group Membership Protocol. Used with multicast protocols to determine whether group members are present.

IGP Interior gateway protocol, such as IS-IS, OSPF, and RIP.

IKE Internet Key Exchange. The key management protocol used in IPSec, IKE combines the ISAKMP and Oakley protocols to create encryption keys and security associations.

import To install routes from the routing protocols into a routing table.

ingress router In MPLS, first router in a label-switched path (LSP). *See also egress router*.

inter-AS routing Routing of packets among different ASs. *See also EBGP*.

intercluster reflection In a BGP route reflection, the redistribution of routing information by a route reflector system to all nonclient peers (BGP peers not in the cluster). *See also route reflection*.

interface routes Routes that are in the routing table because an interface has been configured with an IP address. Also called *direct routes*.

intermediate system In IS-IS, network entity that sends and receives packets and that can also route packets.

internal BGP *See IBGP*.

Internet Key Exchange *See IKE*.

Internet Protocol Security *See IPSec*.

Internet Security Association and Key Management Protocol *See ISAKMP*.

intra-AS routing The routing of packets within a single AS. *See also IBGP*.

IP Internet Protocol. The protocol used for sending data from one point to another on the Internet.

IPSec Internet Protocol Security. The industry standard for establishing VPNs, IPSec comprises a group of protocols and algorithms that provide authentication and encryption of data across IP-based networks.

ISAKMP Internet Security Association and Key Management Protocol. A protocol that allows the receiver of a message to obtain a public key and use digital certificates to authenticate the sender's identity. ISAKMP is designed to be key exchange independent; that is, it supports many different key exchanges. *See also IKE and Oakley*.

IS-IS Intermediate System-to-Intermediate System protocol. Link-state, interior gateway routing protocol for IP networks that also uses the shortest-path first (SPF) algorithm to determine routes.

ISO International Organization for Standardization. Worldwide federation of standards bodies that promotes international standardization and publishes international agreements as International Standards.

ISP Internet service provider. Company that provides access to the Internet and related services.

ITU International Telecommunications Union (formerly known as the CCITT). Group supported by the United Nations that makes recommendations and coordinates the development of telecommunications standards for the entire world.

J

jitter Small random variation introduced into the value of a timer to prevent multiple timer expirations from becoming synchronized.

K

kernel forwarding table *See forwarding table*.

L

label In MPLS, 20-bit unsigned integer in the range 0 through 1048575, used to identify a packet traveling along an LSP.

label-switched path (LSP) Sequence of routers that cooperatively perform MPLS operations for a packet stream. The first router in an LSP is called the *ingress router*, and the last router in the path is called the *egress router*. An LSP is a point-to-point, half-duplex connection from the ingress router to the egress router. (The ingress and egress routers cannot be the same router.)

label switching *See MPLS*.

label-switching router *See LSR*.

link Communication path between two neighbors. A link is *up* when communication is possible between the two end points.

link-state PDU (LSP) Packets that contain information about the state of adjacencies to neighboring systems.

local preference Optional BGP path attribute carried in internal BGP update packets that indicates the degree of preference for an external route.

loose In the context of traffic engineering, a path that can use any route or any number of other intermediate (transit) points to reach the next address in the path. (Definition from RFC 791, modified to fit LSPs.)

LSP *See label-switched path (LSP)) or link-state PDU (LSP).*

LSR Label-switching router. A router on which MPLS and RSVP are enabled and is thus capable of processing label-switched packets.

M

martian address Network address about which all information is ignored.

mask *See subnet mask*.

MBGP Multiprotocol BGP. An extension to BGP that allows you to connect multicast topologies within and between BGP ASs.

MBone Internet multicast backbone. An interconnected set of subnetworks and routers that support the delivery of IP multicast traffic. The MBone is a virtual network that is layered on top of sections of the physical Internet.

MCS Miscellaneous Control Subsystem. On an M160 router, provides control and monitoring functions for router components and SONET clocking for the router.

MD5 Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. It is used in AH and ESP. *See also SHA-1*.

MDRR Modified Deficit Round Robin. A method for selecting queues to be serviced.

MED Multiple exit discriminator. Optional BGP path attribute consisting of a metric value that is used to determine the exit point to a destination when all other factors in determining the exit point are equal.

mesh Network topology in which devices are organized in a manageable, segmented manner with many, often redundant, interconnections between network nodes.

Message Digest 5 *See MD5*.

MIB Management Information Base. Definition of an object that can be managed by SNMP.

midplane Forms the rear of the PIC cage on M5 and M10 routers and the FPC card cage on M20 and M160 routers. Provides data transfer, power distribution, and signal connectivity.

Miscellaneous Control Subsystem *See MCS*.

MPLS Multiprotocol Label Switching. Mechanism for engineering network traffic patterns that functions by assigning to network packets short labels that describe how to forward them through the network. Also called *label switching*. *See also traffic engineering*.

MTBF Mean time between failure. Measure of hardware component reliability.

MTU Maximum transfer unit. Limit on segment size for a network.

multicast Operation of sending network traffic from one network node to multiple network nodes.

multicast distribution tree The data path between the sender (host) and the multicast group member (receiver or listener).

multiprotocol BGP *See MBGP*.

Multiprotocol Label Switching *See MPLS*.

N

neighbor Adjacent system reachable by traversing a single subnetwork. An immediately adjacent router. Also called a *peer*.

NET Network entity title. Network address defined by the ISO network architecture and used in CLNS-based networks.

network layer reachability information *See NLRI*.

network link advertisement An OSPF link-state advertisement flooded throughout a single area by designated routers to describe all routers attached to the network.

Network Time Protocol *See NTP*.

NIC Network Information Center. Internet authority responsible for assigning Internet-related numbers, such as IP addresses and autonomous system numbers. *See also IANA*.

NLRI Network layer reachability information. Information that is carried in BGP packets and is used by MBGP.

nonclient peer In a BGP route reflection, a BGP peer that is not a member of a cluster. *See also client peer*.

not-so-stubby area *See NSSA*.

NSAP Network service access point. Connection to a network that is identified by a network address.

n-selector Last byte of an nonclient peer address.

NSSA Not-so-stubby area. In OSPF, a type of stub area in which external routes can be flooded.

NTP Network Time Protocol. Protocol used to synchronize computer clock times on a network.

O

Oakley A key determination protocol based on the Diffie-Hellman algorithm that provides added security, including authentication. Oakley was the key-exchange algorithm mandated for use with the initial version of ISAKMP, although various algorithms can be used. Oakley describes a series of key exchanges called “modes” and details the services provided by each; for example, Perfect Forward Secrecy for keys, identity protection, and authentication. *See also ISAKMP*.

OC Optical Carrier. In SONET, Optical Carrier levels indicate the transmission rate of digital signals on optical fiber.

OSI Open System Interconnection. Standard reference model for how messages are transmitted between two points on a network.

OSPF Open Shortest Path First. A link-state IGP that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the *Dijkstra algorithm*).

P

package A collection of files that make up a JUNOS software component.

Packet Forwarding Engine The architectural portion of the router that processes packets by forwarding them between input and output interfaces.

path attribute Information about a BGP route, such as the route origin, AS path, and next-hop router.

PCI Peripheral Component Interconnect. Standard, high-speed bus for connecting computer peripherals. Used on the Routing Engine.

PCMCIA Personal Computer Memory Card International Association. Industry group that promotes standards for credit card-size memory or I/O devices.

PDU Protocol data unit. IS-IS packets.

PE router Provider edge router. A router in the service provider's network that is connected to a customer edge (CE) device and that participates in a Virtual Private Network (VPN).

PEC Policing Equivalence Classes. In traffic policing, a set of packets that is treated the same by the packet classifier.

peer An immediately adjacent router with which a protocol relationship has been established. Also called a *neighbor*.

Perfect Forward Secrecy *See PFS*.

PFE *See Packet Forwarding Engine*.

PFS A condition derived from an encryption system that changes encryption keys often and ensures that no two sets of keys have any relation to each other. The advantage of PFS is that if one set of keys is compromised, only communications using those keys are at risk. An example of a system that uses PFS is Diffie-Hellman.

Physical Interface Card *See PIC*.

PIC Physical Interface Card. A network interface–specific card that can be installed on an FPC in the router.

PIM Protocol Independent Multicast. A protocol-independent multicast routing protocol. PIM Sparse Mode routes to multicast groups that might span wide-area and interdomain internets. PIM Dense Mode is a flood-and-prune protocol.

PLP Packet Loss Priority.

PLP bit Packet Loss Priority bit. Used to identify packets that have experienced congestion or are from a transmission that exceeded a service provider's customer service license agreement. This bit can be used as part of a router's congestion control mechanism and can be set by the interface or by a filter.

policing Applying rate limits on bandwidth and burst size for traffic on a particular interface.

pop Removal of the last label, by a router, from a packet as it exits an MPLS domain.

PPP Point-to-Point Protocol. Link-layer protocol that provides multiprotocol encapsulation. It is used for link-layer and network-layer configuration.

precedence bits The first three bits in the ToS byte. On a Juniper Networks router, these bits are used to sort or classify individual packets as they arrive at an interface. The classification determines the queue to which the packet is directed upon transmission.

preference Desirability of a route to become the active route. A route with a lower preference value is more likely to become the active route. The preference is an arbitrary value in the range 0 through 255 that the routing protocol process uses to rank routes received from different protocols, interfaces, or remote systems.

preferred address On an interface, the default local address used for packets sourced by the local router to destinations on the subnet.

primary address On an interface, the address used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface.

primary interface Router interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

Protocol-Independent Multicast	*See PIM*.
provider edge router	*See PE router*.
provider router	Router in the service provider's network that does not attach to a customer edge (CE) device.
PSNP	Partial sequence number PDU. Packet that contains only a partial list of the LSPs in the IS-IS link-state database.
push	Addition of a label or stack of labels, by a router, to a packet as it enters an MPLS domain.

Q

QoS	Quality of service. Performance, such as transmission rates and error rates, of a communications channel or system.
quality of service	*See QoS*.

R

RADIUS	Remote Authentication Dial-In User Service. Authentication method for validating users who attempt to access the router using Telnet.
Random Early Detection	*See RED*.
rate limiting	*See policing*.
RBOC	(Pronounced "are-bock") Regional Bell operating company. Regional telephone companies formed as a result of the divestiture of the Bell System.
RDRAM	RAMBUS dynamic random access memory.
RED	Random Early Detection. Gradual drop profile for a given class that is used for congestion avoidance. RED tries to anticipate incipient congestion and reacts by dropping a small percentage of packets from the head of the queue to ensure that a queue never actually becomes congested.
Rendezvous Point	*See RP*.
Resource Reservation Protocol	*See RSVP*.
RFC	Request for Comments. Internet standard specifications published by the Internet Engineering Task Force.
RFI	Radio frequency interference. Interference from high-frequency electromagnetic waves emanating from electronic devices.
RIP	Routing Information Protocol. Distance-vector interior gateway protocol that makes routing decisions based on hop count.
route flapping	Situation in which BGP systems send an excessive number of update messages to advertise network reachability information.
route identifier	IP address of the router from which a BGP, IGP, or OSPF packet originated.

route reflection In BGP, configuring a group of routers into a cluster and having one system act as a route reflector, redistributing routes from outside the cluster to all routers in the cluster. Routers in a cluster do not need to be fully meshed.

router link advertisement OSPF link-state advertisement flooded throughout a single area by all routers to describe the state and cost of the router's links to the area.

routing domain *See AS*.

Routing Engine Architectural portion of the router that handles all routing protocol processes, as well as other software processes that control the router's interfaces, some of the chassis components, system management, and user access to the router.

routing instance A collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables and the routing protocol parameters control the information in the routing tables.

routing table Common database of routes learned from one or more routing protocols. All routes are maintained by the JUNOS routing protocol process.

RP For PIM-SM, a core router acting as the root of the distribution tree in a shared tree.

rpd JUNOS software routing protocol process (daemon). User-level background process responsible for starting, managing, and stopping the routing protocols on a Juniper Networks router.

RPM Reverse path multicasting. Routing algorithm used by DVMRP to forward multicast traffic.

RSVP Resource Reservation Protocol. Resource reservation setup protocol designed to interact with integrated services on the Internet.

S

SA Security association. An IPSec term that describes an agreement between two parties about what rules to use for authentication and encryption algorithms, key exchange mechanisms, and secure communications.

SAP Session Announcement Protocol. Used with multicast protocols to handle session conference announcements.

SAR Segmentation and reassembly. Buffering used with ATM.

SCB System Control Board. On an M40 router, the part of the Packet Forwarding Engine that performs route lookups, monitors system components, and controls FPC resets.

SCG SONET Clock Generator. Provides Stratum 3 clock signal for the SONET/SDH interfaces on the router. Also provides external clock inputs.

SDH Synchronous Digital Hierarchy. CCITT variation of SONET standard.

SDP Session Description Protocol. Used with multicast protocols to handle session conference announcements.

SDRAM Synchronous dynamic random access memory.

Secure Hash Algorithm *See SHA-1*.

secure shell *See SSH*.

security association *See SA.*

Security Parameter Index *See SPI.*

SFM Switching and Forwarding Module. On an M160 router, a component of the Packet Forwarding Engine that provides route lookup, filtering, and switching to FPCs.

SHA-1 Secure Hash Algorithm. A widely used hash function for use with Digital Signal Standard (DSS). SHA-1 is more secure than MD5.

shortest-path-first algorithm *See SPF.*

signaled path In traffic engineering, an explicit path; that is, a path determined using RSVP signaling. The ERO carried in the packets contains the explicit path information.

SIB Switch Interface Board. Provides the switching function to the destination Packet Forwarding Engine.

simplex interface An interface that assumes that packets it receives from itself are the result of a software loopback process. The interface does not consider these packets when determining whether the interface is functional.

SNMP Simple Network Management Protocol. Protocol governing network management and the monitoring of network devices and their functions.

SONET Synchronous Optical Network. High-speed (up to 2.5 Gbps) synchronous network specification developed by Bellcore and designed to run on optical fiber. STS-1 is the basic building block of SONET. Approved as an international standard in 1988. *See also SDH.*

SPF Shortest-path first, an algorithm used by IS-IS and OSPF to make routing decisions based on the state of network links. Also called the *Dijkstra algorithm*.

SPI Security Parameter Index. A portion of the IPSec Authentication Header that communicates which security protocols, such as authentication and encryption, are used for each packet in a VPN connection.

SPQ Strict Priority Queuing. Dequeuing method that provides a special queue that is serviced until it is empty. The traffic sent to this queue tends to maintain a lower latency and more consistent latency numbers than traffic sent to other queues. *See also APQ.*

SSB System and Switch Board. On an M20 router, Packet Forwarding Engine component that performs route lookups and component monitoring and monitors FPC operation.

SSH Secure shell. Software that provides a secured method of logging in to a remote network system.

SSRAM Synchronous Static Random Access Memory.

static LSP *See static path.*

static path In the context of traffic engineering, a static route that requires hop-by-hop manual configuration. No signaling is used to create or maintain the path. Also called a *static LSP.*

STM Synchronous Transport Module. CCITT specification for SONET at 155.52 Mbps.

strict In the context of traffic engineering, a route that must go directly to the next address in the path. (Definition from RFC 791, modified to fit LSPs.)

STS Synchronous Transport Signal. Synchronous Transport Signal level 1. Basic building block signal of SONET, operating at 51.84 Mbps. Faster SONET rates are defined as STS-*n*, where *n* is a multiple of 51.84 Mbps. *See also SONET*.

stub area In OSPF, an area through which, or into which, AS external advertisements are not flooded.

subnet mask Number of bits of the network address used for host portion of a Class A, Class B, or Class C IP address.

summary link advertisement OSPF link-statement advertisement flooded throughout the advertisement's associated areas by area border routers to describe the routes that they know about in other areas.

sysid System identifier. Portion of the ISO nonclient peer. The sysid can be any 6 bytes that are unique throughout a domain.

System and Switch Board *See SSB*.

T

TACACS + Terminal Access Controller Access Control System Plus. Authentication method for validating users who attempt to access the router using Telnet.

TCP Transmission Control Protocol. Works in conjunction with Internet Protocol (IP) to send data over the Internet. Divides a message into packets and tracks the packets from point of origin to destination.

ToS Type of service. The method of handling traffic using information extracted from the fields in the ToS byte to differentiate packet flows.

traffic engineering Process of selecting the paths chosen by data traffic in order to balance the traffic load on the various links, routers, and switches in the network. (Definition from http://www.ietf.org/internet-drafts/draft-ietf-mpls-framework-04.txt.) *See also MPLS*.

transit area In OSPF, an area used to pass traffic from one adjacent area to the backbone or to another area if the backbone is more than two hops away from an area.

transit router In MPLS, any intermediate router in the LSP between the ingress router and the egress router.

transport mode An IPSec mode of operation in which the data payload is encrypted, but the original IP header is left untouched. The IP addresses of the source or destination can be modified if the packet is intercepted. Because of its construction, transport mode can be used only when the communication endpoint and cryptographic endpoint are the same. VPN gateways that provide encryption and decryption services for protected hosts cannot use transport mode for protected VPN communications. *See also tunnel mode*.

Triple-DES A 168-bit encryption algorithm that encrypts data blocks with three different keys in succession, thus achieving a higher level of encryption. Triple-DES is one of the strongest encryption algorithms available for use in VPNs.

tunnel Private, secure path through an otherwise public network.

tunnel mode An IPSec mode of operation in which the entire IP packet, including the header, is encrypted and authenticated and a new VPN header is added, protecting the entire original packet. This mode can be used by both VPN clients and VPN gateways, and protects communications that come from or go to non-IPSec systems. *See also transport mode*.

Tunnel PIC A physical interface card that allows the router to perform the encapsulation and decapsulation of IP datagrams. The Tunnel PIC supports IP-IP, GRE, and PIM register encapsulation and decapsulation. When the Tunnel PIC is installed, the router can be a PIM rendezvous point (RP) or a PIM first-hop router for a source that is directly connected to the router.

type of service *See ToS*.

U

unicast Operation of sending network traffic from one network node to another individual network node.

UPS Uninterruptible power supply. Device that sits between a power supply and a router (or other piece of equipment) the prevents undesired power-source events, such as outages and surges, from affecting or damaging the device.

V

vapor corrosion inhibitor *See VCI*.

VCI Vapor corrosion inhibitor. Small cylinder packed with the router that prevents corrosion of the chassis and components during shipment.

VCI Virtual circuit identifier. 16-bit field in the header of an ATM cell that indicates the particular virtual circuit the cell takes through a virtual path. Also called a *logical interface. See also VPI*.

virtual circuit identifier *See VCI*.

virtual link In OSPF, a link created between two routers that are part of the backbone but are not physically contiguous.

virtual path identifier *See VPI*.

virtual private network *See VPN*.

Virtual Router Redundancy Protocol *See VRRP*.

VPI virtual path identifier. 8-bit field in the header of an ATM cell that indicates the virtual path the cell takes. *See also VCI*.

VPN virtual private network. A private data network that makes use of a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.

VRRP Virtual Router Redundancy Protocol. On Fast Ethernet and Gigabit Ethernet interfaces, allows you to configure virtual default routers.

W

wavelength-division multiplexing *See WDM.*

WDM Wavelength-division multiplexing. Technique for transmitting a mix of voice, data, and video over various wavelengths (colors) of light.

WFQ Weighted Fair Queuing.

weighted round-robin *See WRR.*

WRR Weighted round-robin. Scheme used to decide the queue from which the next packet should be transmitted.

Part 10
Index

- Index on page 849
- Index of Statements and Commands on page 861

Index

Index

Bold numbers point to command summary pages.

Symbols

!
    wildcard character....188
#, in configuration statements.... xxxix
$, regular expression operator....13
( )
    in syntax descriptions xxxviii
    regular expression operator....13
*
    wildcard character....188
-, wildcard character....188
< >, in syntax descriptions.... xxxviii
>
    prompt....3
[ ]
    in configuration statements.... xxxviii
    regular expression operator....13
    wildcard character....188
^
    CLI error messages....9
    regular expression operator....13
{ }, in configuration statements.... xxxviii
|
    in syntax descriptions xxxviii
| (pipe) command....22, 22

A

accounting options
    profile 178
    records....178
active prunes....657
addresses
    multicast announcements 688
    routing table entries 555, 565
administrative groups, MPLS 696
administratively scoped multicast information 671
advertisement of routing information....527
alarm conditions....135
/altconfig file system....57
/altroot file system....57
announced multicast sessions 672
ARP table....44
AS paths 514
ATM interface 209
ATM node reachability....50

B

backing up partitions 57
BGP
    damping parameters 613
    monitoring 613
    neighbors 614
boot messages....86
bootstrap router....676
braces, in configuration statements xxxviii
brackets
    angle, in syntax descriptions xxxviii
    square, in configuration statements xxxviii
buffers, system 88
byte counts, firewall filters....475

C

CB....171
CCC
    connections....693
Channelized DS-3 to DS-0 interfaces....225
Channelized DS-3 to DS-1 interfaces....239
Channelized E1 interfaces 253
Channelized OC-12 interfaces....269
Channelized STM-1 interfaces....283
chassis
    alarm conditions 135
    FEB....150
    firmware version 150
    FPCs....151, 153
    hardware installed....153
    monitoring commands....127, 167
    PICs....151, 153
    See also Routing Engines

chassis environment 169
checksum values 85
class-of-service
monitoring commands..... 487
clear arp command 44
clear arp vpn command 44
clear bgp damping command 613
clear bgp neighbor command 614
clear command..... 20
clear firewall command 475
clear helper statistics command 479
clear igmp membership command..... 641
clear igmp statistics command..... 643
clear ike security-associations command 817
clear interfaces statistics command 193
clear ipsec security-associations command 818
clear IPv6 neighbors command..... 635
clear ipv6 router-advertisement command..... 636
clear isis adjacency command..... 574
clear isis database command 575
clear isis statistics command..... 576
clear ldp neighbor command 737
clear ldp statistics command..... 737
clear log command 44
clear mpls lsp command..... 692
clear msdp cache command 644
clear msdp statistics command 644
clear multicast statistics command 644
clear ospf database command..... 594
clear ospf neighbor command..... 594
clear ospf statistics command 595
clear pim join command..... 645
clear pim statistics command..... 646
clear rip statistics command 625
clear ripng statistics command..... 629
clear route forwarding-table command 513
clear rsvp session command..... 717
clear rsvp statistics command 718
clear snmp statistics command 118
clear system reboot command 44
clear vrrp command..... 465
CLI
command completion..... 4, 8, 17, 25
command hierarchy..... 4
command history..... 19, 29, 66
editing command line 10
environment commands..... 25
environment settings 16
exiting..... 51
help about commands 4, 6
messages 9
monitoring system usage 19
operational commands 20
overview 3
command completion 4, 8, 17, 25
command hierarchy..... 4
command history..... 19, 29, 66
command output
counting lines..... 15
filtering 13, 16, 22
output fields
See also descriptions of specific commands
retaining..... 15
saving to files 12
screen display 10, 15
scrolling 11
searching for strings..... 13
system IDs in IS-IS command output 573
command-line interface *See* CLI
commands
command completion 4, 8, 17, 25
command hierarchy..... 4
filenames, specifying..... 43
help about commands..... 4, 6
output *See* command output
router management commands
See commands for router management
URLs, specifying..... 43
commands for router management
BERT interfaces..... 459
BGP 613
chassis 127, 167
class-of-service 487
environment commands, CLI..... 25
firewall filters 475
interfaces 187, 445
IP multicast 639
IS-IS..... 573
Layer2 VPN 751
LDP 737
MPLS..... 691, 817
operational..... 20, 36
OSPF 593
protocol-independent routing properties 512
RIP 625, 629
routing policy 483
RSVP 717
SNMP 117
system management..... 41
troubleshooting..... 36
VPN..... 779
VRRP interfaces 465
comments, in configuration statements xxxix
completion of commands..... 4, 8, 17, 25
/config file system 57
configuration files, verifying syntax..... 113
configure command 20
configure exclusive command 44
connections
CCC connections 693
IP socket information 89
secure shell connections..... 110, 812
connectivity, testing 48
constrained shortest path first statistics..... 697

contents of files 46
Control Board 171
conventions, documentation xxxvii
CoS
monitoring 487
count command, usage guidelines 15
counting lines in command output 15
craft interface
troubleshooting routers 37
CSPF statistics 697
curly braces, in configuration statements xxxviii
current time 101
customer support
contacting xl
help about commands 6

D

damped routes 613
date
setting 18, 27
setting from NTP servers 18, 28
deleting files 45
device flags output field 190
Direct Memory Access statistics 79
disk space available 100
DMA statistics 79
DNS host names 68
documentation conventions xxxvii
DVMRP
active groups, usage information 675
active prunes 657
interface information 653
neighbors 654
prefixes 655
tunnels 674

E

E1 interfaces 299
E3 interfaces 313
Emacs keyboard sequences 10
enabled output field 190
entries in routing tables *See* route entries
environment commands, CLI 25
environment, CLI
configuring 16
date setting 27
date setting from NTP servers 28
displaying environment settings 18, 28
example settings 18
idle timeout setting 17, 25
prompt 17, 26
restarting after software upgrade 17
restarting software upgrade 26
screen dimensions 16, 26, 27
terminal type 16, 27
time setting 27
time setting from NTP servers 28
error messages, CLI 9
errors, Packet Forwarding Engine 80
ES interfaces 325
except command 13
exiting CLI 51
extensions, loaded 97

F

family flags output field 192
Fast Ethernet interfaces 335
FEB 150
fields, output *See* output fields
file command 20
file copy command 45
file delete command 45
file list command 46
file rename command 46
file show command 46
file systems
free disk space 100
partitions, backing up 57
state of files 85
See also files
filenames, specifying in commands 43
files
checksum values 85
contents, displaying 46
copying 45
deleting 45
listing 46
renaming 46
saving command output 12
See also file systems
filtering command output 13, 16, 22
filters *See* firewall filters
firewall filters
monitoring commands 475
packet and byte counts 475
firmware version level 150
Forwarding Engine Board 150
forwarding tables
multicast forwarding table 666
FPC statistics 72
FPCs
firmware version level 150
status information 151
free disk space 100

G
Gigabit Ethernet interfaces 347
GRE interfaces 361
groups
DVMRP 675
IGMP 658
MPLS 696
PIM 675, 679

H
halts
clearing 44
displaying pending halts 97
halting router software 53, 97
hardware installed, listing 153
help about commands 4, 6
help command 20
hierarchy of commands 4
history, CLI commands 19, 29, 66
hold command, usage guidelines 15
host names, DNS 68
host reachability, testing 48

I
idle timeout setting, CLI environment 17, 25
IGMP
group membership 658
interface information 660
interface flags output field 191
interfaces
BERT monitoring 459
monitor traffic match conditions 445
monitoring 187, 445
VRRP monitoring 465
wildcard characters, in interface names 187
See also craft interface
See also names of specific interfaces
interface-specific filter
status 476
internal NTP variables 71
IP multicast
addresses listened to 688
administratively scoped multicast information 671
announced multicast sessions 672
DVMRP
active groups 675
active prunes 657
interface information 653
neighbors 654
prefixes 655
tunnels 674
IGMP
group membership 658
interface information 660
monitoring 639
MSDP
peers 663
statistics 663
multicast forwarding table 666
multicast next-hop table 665
multicast reverse-path-forwarding calculations 669
multicast router information 664
multicast statistics 673
PIM
active groups 675
bootstrap routers 676
groups 679
interface information 677
neighbors 682
RPs 683
source RPF state 686
statistics 687
tunnels 674
IP socket information 89
IP statistics 81
IP-IP interfaces 371
IPSEC
clear ike security-associations command 817
clear ipsec security-associations command 818
show ipsec security-associations command 819, 822
IPv6 632
clear neighbors 635
clear router-advertisement 636
monitoring 635
show ipv6 neighbors 636
show route protocol ripng 631
show router-advertisement 637
IS-IS
monitoring 573
system ID format 573
isolating cause of network problems 32, 34

K
keyboard sequences 10

L

label-switched paths *See* LSPs
Layer 2 VPNs
monitoring 751
show bgp neighbor command 752
show bgp summary command 757
show l2circuit connections 760
show l2vpn connections command 765
show route advertising-protocol command 769
show route command 768
show route forwarding-table command 771
show route receive-protocol command 774
show route table command 777
Layer 3 VPNs
reachability 780
show bgp neighbor command 781
show bgp summary command 787
show ospf database command 789
show ospf interface command 790
show ospf log command 790
show ospf neighbor command 791
show ospf route command 791
show ospf statistics command 792
show rip neighbor command 793
show rip statistics command 793
show route advertising-protocol command 795
show route command 794
show route forwarding-table command 797
show route instance command 800
show route protocol command 802
show route receive-protocol command 805
show route table command 809
ssh command 812
telnet command 813
traceroute command 814
LDP
monitoring 737
link flags output field 191
listing files 46
loads *See* environmental information
log files
displaying contents 69
listing 69
monitoring entries 47, 452
removing contents 44
logging
system operations 19
logical interface flags output field 192

M

mastership election, Routing Engines 130
match command 13
messages
CLI messages 9
displaying to users 52
system boot messages 86
monitor command 21, 47, 452
monitor interface command 449
monitor traffic command 452
monitor traffic match conditions 445
MPLS
administrative groups 696
CCC connections 693
CSPF statistics 697
monitoring 691, 817
MSDP
peers 663
statistics 663
mtrace command 21, 647
mtrace from-source command 648
mtrace monitor command 650
mtrace to-gateway command 651
multicast
forwarding table 666
next-hop table 665
reverse-path-forwarding calculations 669
See IP multicast
multicast sessions, announced 672
multicast statistics 673
Multilink Services interfaces 381

N

names
filenames 43
wildcard characters in interface names 187
neighbors
BGP 614
DVMRP 654
ipv6 635
PIM 682
network connectivity, testing 48
networks, troubleshooting 36
strategies 31, 34
tools 37
next-hops
multicast next-hop table 665
PFE next-hops 73
route entries sent to 555
no-more command 15
notification statistics 82
NTP
internal variables 71
NTP peers 70

O

operational mode, CLI
- commands 20, 36
- troubleshooting routers 36

OSPF
- monitoring 593

output fields
- show interfaces commands
  - device flags field 190
  - enabled field 190
  - family flags field 192
  - interface flags field 191
  - link flags field 191
  - logical interface flags field 192
- *See also* command output
- *See also descriptions of specific commands*

P

packages
- rollback from router 62
- upgrading 57

packet counts, firewall filters 475

Packet Forwarding Engine
- DMA statistics 79
- error statistics 80
- FPC statistics 72
- IP statistics 81
- next hops 73
- notification statistics 82
- polled I/O statistics 82
- routing table 74
- SCB information 75
- SFM information 77
- terse information display 84
- traffic statistics 83

partitions, backing up 57

PEM 172

permissions, displaying 65

PICs 151

PIM
- bootstrap router 676
- groups 675, 679
- interface information 677
- neighbors 682
- RPs 683
- source RPF state 686
- statistics 687
- tunnels 674

ping atm command 50

ping command 21, 48

ping vpn-interface command 51, 780

policies *See* routing policies

polled I/O statistics 82

Power Entry Module 172

prefixes
- DVMRP 655

processes
- displaying information 91
- restarting 64

prompt, CLI 3, 17, 26

protocols
- protocol-related statistics 98
- system uptime 101
- *See also names of specific protocols*

prunes, DVMRP 657

Q

quit command 21, 51

R

reachability, testing 48, 50, 51, 780

rebooting router software
- clearing reboots 44
- displaying pending reboots 97
- rebooting 55

redirecting command output 12

remote systems
- secure shell connections 110, 812

renaming files 46

rendezvous points, PIM 683

request chassis fpc slot command 128

request chassis pcg slot command 128

request chassis pic command 129

request chassis routing-engine master command 130

request chassis scg command 167

request chassis sfm master switch command 132

request chassis sfm slot command 131

request chassis sib command 168

request chassis spmb restart slot command 168

request chassis ssb master switch command 133

request command 22

request message command 52

request support information command 52

request system halt command 53

request system logout user command 54

request system reboot command 55

request system snapshot command 57

request system software add command 59

request system software command 63

request system software delete command 61

request system software rollback command 62

restart command 23, 64

restart helper statistics 481

restarting after software upgrade 17, 26

restarting software processes 64

retaining command output 15

reverse-path-forwarding calculations 669

RIP
- monitoring 625, 629

rollback software from router 62

route advertisement 527
route entries
addresses in 555, 565
learned from specified addresses 565
sent to next-hop addresses 555
See also routing tables
route instance
route instance 550
router chassis *See* chassis
router management commands
See commands for router management
router software
halting 53, 97
loaded extensions 97
rebooting 55, 97
rollback from router 62
system uptime 101
upgrading 59, 63
routers
bootstrap router 676
chassis *See* chassis
removing router software 63
rollback router software 62
router management commands
See commands for router management
software *See* router software
system uptime 101
troubleshooting 31, 34, 36, 37
routes to hosts, determining 814
Routing Engines
forwarding tables 666
mastership election 130
routing tables
entries *See* route entries
Packet Forwarding Engine 74
route advertisement 527
RPs 683
RSVP
monitoring 717

S

save command 12
save command, usage guidelines 12
saving command output 15
SCB
displaying information 75
firmware version level 150
SCG 167, 173
screen display
command output 10, 15
screen dimensions 16, 26, 27
searching command output 13
secure shell connections 110, 812
set chassis display message command 134
set cli complete-on-space command 25
usage guidelines 17
set cli idle-timeout command 25
usage guidelines 17
set cli prompt command 26
usage guidelines 17
set cli restart-on-upgrade command 26
usage guidelines 17
set cli screen-length command 26
usage guidelines 16
set cli screen-width command 27
usage guidelines 16
set cli terminal command 27
usage guidelines 16
set command 23
usage guidelines 16
set date command 27
usage guidelines 18
set date ntp command 28
usage guidelines 18
SFMs
firmware version level 150
status information 77
show accounting profile command 178
show accounting records command 182
show aps command 455
show arp command 65
show as-path command 514
show bgp group command 615
show bgp neighbor command 617, 752, 781
show bgp summary command 622, 757, 787
show chassis alarm command 135
show chassis craft-interface command 136
show chassis environment cb command 171
show chassis environment command 138, 169
show chassis environment fpc command 141
show chassis environment fpm command 142
show chassis environment mcs command 143
show chassis environment pcg command 144
show chassis environment pem command 145, 172
show chassis environment routing engine
command 146
show chassis environment scg command 173
show chassis environment sfm command 147
show chassis environment sib command 174
show chassis feb command 149
show chassis firmware command 150
show chassis fpc command 151
show chassis hardware command 153
show chassis mac-addresses command 157
show chassis routing-engine command 158
show chassis scb command 162
show chassis sfm command 163
show chassis spmb command 174
show chassis spmb sibs command 175
show chassis ssb command 165
show class-of-service classifier command 489
show class-of-service code-point-alias command 490
show class-of-service command 488

show class-of-service drop-profile command 491
show class-of-service forwarding-class command 492
show class-of-service forwarding-table classifier command 493
show class-of-service forwarding-table classifier mapping command 494
show class-of-service forwarding-table command 492
show class-of-service forwarding-table drop-profile command 495
show class-of-service forwarding-table rewrite-rule command 496
show class-of-service forwarding-table rewrite-rule mapping command 497
show class-of-service forwarding-table scheduler-map command 498
show class-of-service interface command 500
show class-of-service rewrite-rule command 501
show class-of-service scheduler-map command 502
show cli authorization command 65
show cli command 28
usage guidelines 18
show cli history command 29, 66
usage guidelines 19
show command 23
show configuration command 67
show connections command 693
show dvmrp interfaces command 653
show dvmrp neighbors command 654
show dvmrp prefix command 655
show dvmrp prunes command 657
show firewall filter command 476
show firewall log command 477
show helper statistics command 480
show host command 68
show igmp group command 658
show igmp interface command 660
show igmp statistics command 661
show interface extensive command
CoS related attributes only 504
show interfaces commands 194
ATM interfaces 209
Channelized DS-3 to DS-0 interfaces 225
Channelized DS-3 to DS-1 interfaces 239
Channelized E1 interfaces 253
Channelized OC-12 interfaces 269
Channelized STM-1 interfaces 283
E1 interfaces 299
E3 interfaces 313
ES interfaces 325
Fast Ethernet interfaces 335
Gigabit Ethernet interfaces 347
GRE interfaces 361
IP-IP interfaces 371
Multilink Services interfaces 381
SONET/SDH interfaces 393
T1 interfaces 409
T3 interfaces 421
Virtual Tunnel (VT) 435
show interfaces descriptions command 195
show interfaces filters | policers command 196
show interfaces media commands 199
show interfaces queue command 506
show interfaces routing command 200
show interfaces statistics commands 205
show interfaces terse command 206, 759
show ipsec security-associations command 819, 822
show ipv6 neighbors command 631, 636
show isis adjacency command 577
show isis database command 580
show isis hostname command 583
show isis interface command 584
show isis routes command 587
show isis spf command 588
show isis statistics command 591
show l2circuit connections 760
show l2vpn connections command 765
show ldp database command 738
show ldp interface command 741
show ldp neighbor command 742
show ldp path command 743
show ldp route command 745
show ldp session command 747
show ldp statistics command 749
show link-management command 695
show log command 69
show mpls admin-groups command 696
show mpls cspf command 697
show mpls interface command 698
show mpls lsp command 699
show mpls path command 709
show msdp command 662
show msdp statistics command 663
show multicast mrinfo command 664
show multicast nexthops command 665
show multicast route command 666
show multicast rpf command 669
show multicast scope command 671
show multicast sessions command 672
show multicast statistics command 673
show multicast tunnels command 674
show multicast usage command 675
show ntp associations command 70
show ntp status command 71
show ospf database command 596, 789
show ospf interface command 602, 790
show ospf io-statistics command 605
show ospf log command 606, 790
show ospf neighbor command 607, 791
show ospf route command 610, 791
show ospf statistics command 612, 792
show pfe fpc command 72
show pfe next-hop command 73

show pfe route command 74
show pfe scb command 75
show pfe sfm command 77
show pfe statistics dma command 79
show pfe statistics error command 80
show pfe statistics ip command 81
show pfe statistics notification command 82
show pfe statistics pio command 82
show pfe statistics traffic command 83
show pfe terse command 84
show pim bootstrap command 676
show pim interfaces command 677
show pim join command 679
show pim neighbors command 682
show pim rps command 683
show pim source command 686
show pim statistics command 687
show policy command 483
show policy damping command 484
show rip neighbor command 626, 793
show rip statistics command 627, 793
show ripng neighbor command 630
show ripng statistics 632
show ripng statistics command 632
show route advertising-protocol
command 527, 769, 795
show route aspath-regex command 529
show route best command 531
show route command 516, 768, 794
show route community command 532
show route damping command 534
show route exact command 537
show route export command 539
show route export vrf-target command 541
show route forwarding-table command 543, 771, 797
show route inactive command 549
show route instance command 550, 800
show route label-switched-path command 552
show route martians command 554
show route next-hop command 555
show route output command 556
show route protocol command 557, 802
show route range command 559
show route receive-protocol command 561, 774, 805
show route resolution command 563
show route source-gateway command 565
show route summary command 566
show route table command 567, 777, 809
show route terse command 571
show router-advertisement ipv6 command 637
show rsvp interface command 719
show rsvp neighbor command 724
show rsvp session command 727
show rsvp statistics command 734
show rsvp version command 736
show sap listen command 688
show snmp rmon command 119
show snmp statistics command 123
show system audit command 85
show system boot-messages 86
show system buffers command 88
show system connections 89
show system processes command 91
show system reboot command 97
show system software command 97
show system statistics command 98
show system storage command 100
show system uptime command 101
show system users command 102
show system virtual-memory command 103
show task command 107
show ted database command 710
show ted link command 714
show ted protocol command 716
show version command 109
show vrrp command 466
SIB 174
request 168
SNMP
show snmp rmon command 119
SNMP, monitoring 117
sockets 89
software processes
displaying information 91
restarting 64
software reboots
clearing 44
pending reboots 97
software, router
halting 53, 97
loaded extensions 97
rebooting 44, 55, 97
removing from router 63
rollback from router 62
system uptime 101
upgrading 59, 63
Sonet Clock Generator 167, 173
SONET/SDH interfaces 393
source RPF state, PIM 686
space character, command completion 17, 25
SPMB 174, 175
restart slot 168
SSB 150
ssh command 24, 110, 812
start command 24, 111
statistics
real-time display 449
strategies, troubleshooting 31, 34
strings, searching command output 13
support, technical
customer support, contacting xl
Switch Interface Board 174
request 168
Switch Processor Mezzanine Board 174, 175

Switching and Forwarding Modules 77
symptoms of network problems, identifying 32, 34
syntax of configuration files, verifying 113
syslog statement, usage guidelines 19
system boot messages 86
system buffer information 88
system IDs, IS-IS command output 573
system management commands 41
system operations, logging 19
system usage, monitoring 19

T

T1 interfaces 409
T3 interfaces 421
tab character, command completion 17, 25
tcpdump *See monitor traffic*
technical support
    customer support, contacting xl
    help about commands 4, 6
    system information, displaying 52
telnet command 24, 112, 813
temperature *See* environmental information
terminal type 16, 27
test command 24
test configuration command 113
test interface bert-start command 459
test interface bert-stop command 460
test interface fdl-line-loop command 461
test interface fdl-payload-loop command 462
test interface feac-loop-initiate command 463
test interface feac-loop-terminate command 463
test msdp command 688
test policy command 484
time
    current time 101
    setting 18, 27
    setting from NTP servers 18, 28
    system uptime 101
trace files, monitoring entries 47, 452
traceroute command 24, 114, 814
traffic
    statistics 83
    troubleshooting 31, 34, 36, 37
Traffic Engineering Database *See* TED
troubleshooting 36, 37
    strategies 31, 34
tunnels
    DVMRP 674
    PIM 674
typefaces, documentation conventions xxxvii

U

upgrading software
    installing packages 59, 63
    restarting after upgrades 17, 26
uptime 101
URLs, specifying in commands 43
users
    current user's permissions 65
    logged into router 69

V

verifying syntax of configuration files 113
version information
    firmware 150
Virtual Tunnel (VT) interfaces 435
voltage *See* environmental information
VPNs
    clear arp 44
    monitoring 779
    reachability 51
    routes to hosts, determining 114
    ssh 110
    traceroute 114
VPNs, Layer 2
    advertising-protocol command 769
    show bgp neighbor command 752
    show bgp summary command 757
    show l2circuit connections 760
    show l2vpn connections command 765
    show route command 768
    show route forwarding-table command 771
    show route receive-protocol command 774
    show route table command 777
VPNs, Layer 3
    reachability 780
    show bgp neighbor command 781
    show bgp summary command 787
    show ospf interface command 790
    show ospf log command 790
    show ospf neighbor command 791
    show ospf route command 791
    show ospf statistics command 792
    show route advertising-protocol command 795
    show route command 794
    show route forwarding-table command 797
    show route instance command 800
    show route receive-protocol command 805
    show route table command 809
    show show ospf database command 789
    show show rip neighbor command 793
    show show rip statistics command 793
    show show route protocol command 802
    ssh command 812
    telnet command 813
    traceroute command 814

VRRP
statistics 465

W

wildcard characters
! 188
* 188
in interface names 187

Index

Index of Statements and Commands

Symbols

| (pipe) command 22

C

clear arp command 44
clear bgp damping command 613
clear bgp neighbor command 614
clear command 20
clear firewall command 475
clear helper statistics command 479
clear igmp membership command 641
clear igmp statistics command 643
clear ike security-associations command 817
clear interfaces statistics command 193
clear ipsec security-associations command 818
clear IPv6 neighbors command 635
clear ipv6 router-advertisement command 636
clear isis adjacency command 574
clear isis database command 575
clear isis statistics command 576
clear ldp neighbor command 737
clear ldp statistics command 737
clear log command 44
clear mpls lsp command 692
clear msdp cache command 644
clear msdp statistics command 644
clear multicast statistics command 644
clear ospf database command 594
clear ospf neighbor command 594
clear ospf statistics command 595
clear pim join command 645
clear pim statistics command 646
clear rip statistics command 625
clear ripng statistics command 629
clear route forwarding-table command 513
clear rsvp session command 717
clear rsvp statistics command 718
clear snmp statistics command 118
clear system reboot command 44
clear vrrp command 465
configure command 20
configure exclusive command 44

E

except command 13

F

file command 20
file copy command 45
file delete command 45
file list command 46
file rename command 46
file show command 46

H

help command 20

M

match command 13
monitor command 21, 47, 452
monitor interface command 449
monitor traffic command 452
mtrace command 21, 647
mtrace from-source command 648
mtrace monitor command 650
mtrace to-gateway command 651

N

no-more command 15

P

ping atm command 50
ping command 21, 48
ping vpn-interface command 51, 780

Q

quit command 21, 51

R

request chassis fpc slot command 128
request chassis pcg slot command..... 128
request chassis pic command 129
request chassis routing-engine master
command 130
request chassis scg command 167
request chassis sfm master switch command..... 132
request chassis sfm slot command 131
request chassis sib command..... 168
request chassis spmb restart slot command 168
request chassis ssb master switch command..... 133
request command 22
request message command 52
request support information command..... 52
request system halt command 53
request system logout user command 54
request system reboot command 55
request system snapshot command 57
request system software add command..... 59
request system software command 63
request system software delete command 61
request system software rollback command..... 62
restart command 23, 64
restart helper statistics..... 481

S

save command 12
set chassis display message command 134
set cli complete-on-space command 25
set cli idle-timeout command 25
set cli prompt command..... 26
set cli restart-on-upgrade command 26
set cli screen-length command 26
set cli screen-width command..... 27
set cli terminal command 27
set command..... 23
set date command 27
set date ntp command 28
show accounting profile command..... 178
show accounting records command 182
show aps command 455
show arp command 65
show as-path command..... 514
show bgp group command 615
show bgp neighbor command..... 617, 752, 781
show bgp summary command..... 622, 757, 787
show chassis alarm command 135
show chassis craft-interface command 136
show chassis environment cb command..... 171
show chassis environment command..... 138, 169
show chassis environment fpc command 141
show chassis environment fpm command 142
show chassis environment mcs command 143
show chassis environment pcg command..... 144
show chassis environment pem
command 145, 172
show chassis environment routing engine
command 146
show chassis environment scg command 173
show chassis environment sfm command..... 147
show chassis environment sib command 174
show chassis feb command 149
show chassis firmware command..... 150
show chassis fpc command 151
show chassis hardware command 153
show chassis mac-addresses command 157
show chassis routing-engine command..... 158
show chassis scb command 162
show chassis sfm command 163
show chassis spmb command..... 174
show chassis ssb command 165
show class-of-service classifier command..... 489
show class-of-service code-point-alias
command 490
show class-of-service command 488
show class-of-service drop-profile command 491
show class-of-service forwarding-class
command 492
show class-of-service forwarding-table classifier
command 493
show class-of-service forwarding-table classifier
mapping command..... 494
show class-of-service forwarding-table
command 492
show class-of-service forwarding-table
drop-profile command 495
show class-of-service forwarding-table
rewrite-rule command..... 496
show class-of-service forwarding-table
rewrite-rule mapping command..... 497
show class-of-service forwarding-table
scheduler-map command 498
show class-of-service interface command..... 500
show class-of-service rewrite-rule command..... 501
show class-of-service scheduler-map
command 502
show cli authorization command..... 65
show cli command 28
show cli history command 66
show command 23
show configuration command..... 67
show connections command..... 693
show dvmrp interfaces command 653
show dvmrp neighbors command 654
show dvmrp prefix command 655
show dvmrp prunes command..... 657
show firewall filter command..... 476
show firewall log command 477
show helper statistics command..... 480
show host command 68
show igmp group command 658

show igmp interface command 660
show igmp statistics command 661
show interface extensive command
CoS related attributes only 504
show interfaces commands 194
ATM interfaces 209
Channelized DS-3 to DS-0 interfaces 225
Channelized DS-3 to DS-1 interfaces 239
Channelized E1 interfaces 253
Channelized OC-12 interfaces 269
Channelized STM-1 interfaces 283
E1 interfaces 299
E3 interfaces 313
ES interfaces 325
Fast Ethernet interfaces 335
Gigabit Ethernet interfaces 347
GRE interfaces 361
IP-IP interfaces 371
Multilink Services interfaces 381
SONET/SDH interfaces 393
T1 interfaces 409
T3 interfaces 421
Virtual Tunnel (VT) 435
show interfaces descriptions command 195
show interfaces filters | policers command 196
show interfaces media commands 199
show interfaces queue command 506
show interfaces routing command 200
show interfaces statistics commands 205
show interfaces terse command 206, 759
show ipsec security-associations
command 819, 822
show ipv6 neighbors command 631, 636
show isis adjacency command 577
show isis database command 580
show isis hostname command 583
show isis interface command 584
show isis routes command 587
show isis spf command 588
show isis statistics command 591
show l2circuit connections 760
show l2vpn connections command 765
show ldp database command 738
show ldp interface command 741
show ldp neighbor command 742
show ldp path command 743
show ldp route command 745
show ldp session command 747
show ldp statistics command 749
show link-management command 695
show log command 69
show mpls admin-groups command 696
show mpls cspf command 697
show mpls interface command 698
show mpls lsp command 699
show mpls path command 709
show msdp command 662
show msdp statistics command 663
show multicast mrinfo command 664
show multicast nexthops command 665
show multicast route command 666
show multicast rpf command 669
show multicast scope command 671
show multicast sessions command 672
show multicast statistics command 673
show multicast tunnels command 674
show multicast usage command 675
show ntp associations command 70
show ntp status command 71
show ospf database command 596, 789
show ospf interface command 602, 790
show ospf io-statistics command 605
show ospf log command 606, 790
show ospf neighbor command 607, 791
show ospf route command 610, 791
show ospf statistics command 612, 792
show pfe fpc command 72
show pfe next-hop command 73
show pfe route command 74
show pfe scb command 75
show pfe sfm command 77
show pfe statistics dma command 79
show pfe statistics error command 80
show pfe statistics ip command 81
show pfe statistics notification command 82
show pfe statistics pio command 82
show pfe statistics traffic command 83
show pfe terse command 84
show pim bootstrap command 676
show pim interfaces command 677
show pim join command 679
show pim neighbors command 682
show pim rps command 683
show pim source command 686
show pim statistics command 687
show policy command 483
show policy damping command 484
show rip neighbor command 626, 793
show rip statistics command 627, 793
show ripng neighbor command 630
show ripng statistics command 632
show route advertising-protocol
command 527, 769, 795
show route aspath-regex command 529
show route best command 531
show route command 516, 768, 794
show route community command 532
show route damping command 534
show route exact command 537
show route export command 539
show route export vrf-target command 541
show route forwarding-table
command 543, 771, 797
show route inactive command 549

show route instance command 550, 800
show route label-switched-path command 552
show route martians command..... 554
show route next-hop command..... 555
show route output command 556
show route protocol command..... 557, 802
show route range command 559
show route receive-protocol
command 561, 774, 805
show route resolution command 563
show route source-gateway command 565
show route summary command..... 566
show route table command 567, 777, 809
show route terse command 571
show rsvp interface command 719
show rsvp neighbor command 724
show rsvp session command 727
show rsvp statistics command 734
show rsvp version command 736
show sap listen command 688
show snmp rmon command..... 119
show snmp statistics command 123
show system audit command 85
show system boot-messages 86
show system buffers command..... 88
show system connections 89
show system processes command..... 91
show system reboot command 97
show system software command 97
show system statistics command 98
show system storage command 100
show system uptime command 101
show system users command 102
show system virtual-memory command 103
show task command 107
show ted database command..... 710
show ted link command 714
show ted protocol command..... 716
show version command 109
show vrrp command 466
ssh command 24, 110, 812
start command 24, 111

T

telnet command 24, 112, 813
test command 24
test configuration command 113
test interface bert-start command 459
test interface bert-stop command..... 460
test interface fdl-line-loop command 461
test interface fdl-payload-loop command 462
test interface feac-loop-initiate command 463
test interface feac-loop-terminate command 463
test msdp command 688
test policy command..... 484
traceroute command..... 24, 114, 814